Under consideration for publication in Formal Aspects of Computing

arXiv:1609.00195v1 [cs.LO] 1 Sep 2016

Designing a semantic model for a wide-spectrum language with concurrency

Robert J. Colvin

Ian J. Hayes

Larissa A. Meinicke

School of Information Technology and Electrical Engineering,
The University of Queensland, Australia

Abstract. A wide-spectrum language integrates specification constructs into a programming language in a manner that treats a specification command just like any other command. This paper investigates a semantic model for a wide-spectrum language that supports concurrency and a refinement calculus. In order to handle specifications with rely and guarantee conditions, the model includes explicit environment steps as well as program steps. A novelty of our approach is that we define a set of primitive commands and operators, from which more complex specification and programming language commands are built. The primitives have simple algebraic properties which support proof using algebraic reasoning. The model is general enough to specify notions as diverse as rely-guarantee reasoning, temporal logic, and progress properties of programs, and supports refining specifications to code. It also forms an instance of an abstract concurrent program algebra, which facilitates reasoning about properties of the model at a high level of abstraction.

Keywords: Refinement calculus; wide-spectrum language; concurrency; program algebra; rely-guarantee

1. Introduction and motivation

Concurrent programming languages and their semantics are relatively well understood but the problem of how to abstractly specify concurrent programs remains a challenge. For a sequential program one way to abstractly specify its behaviour is via a precondition-postcondition pair [11, 20]. For concurrent programs, attending only to initial and final states is inadequate: one needs to express the intermediate behaviour of a program and its interaction with its environment along with any assumptions it makes about its environment.

Correspondence and offprint requests to: R. Colvin, I. J. Hayes and L. Meinicke (2 September 2016)

1.1. Goals

The goal of this research is to provide a semantic model for a *concurrent wide-spectrum language* that includes both specification constructs and standard programming combinators and primitives. The benefit lies in expressing related concepts from concurrency in a uniform manner, and thus supporting reuse of theories.

The model supports a *refinement calculus* [2] with a refinement relation, $c_1 \sqsubseteq c_2$, representing that specification c_1 is refined (or implemented) by c_2. The refinement calculus approach is compositional in the sense that refining a subcomponent refines the whole, for example, if $c_1 \sqsubseteq c_2$ and $d_1 \sqsubseteq d_2$ then the parallel composition of c_1 and d_1, $c_1 \parallel d_1$, is refined by $c_2 \parallel d_2$.[1] Extending the refinement calculus to handle concurrency in a compositional manner requires a rich specification language that can state the desired behaviour of a program with respect to assumptions about the behaviour of its environment. The underlying model of a command is based on traces that distinguish program and environment steps originally suggested by Aczel [1] and further developed by de Roever and others [7, 8]. The approach uses a small set of primitive commands and basic operators to define more complex commands.

The original motivation for this work was to provide a general and abstract theory to support the *rely-guarantee* approach for reasoning about concurrent programs of Jones [22, 23] but the resulting theory is applicable to other concepts including atomic operations and temporal logic specifications [27]. The model is an instance of a concurrent program algebra [13], and so it can be used to prove consistency of the axioms given there. The algebra can then be used to verify properties of the model at a higher level of abstraction. The algebra of Hayes [13] is similar to the Concurrent Kleene Algebra of Hoare et al. [21], although they are distinguished by their subtly different axiomatizations. In general, program algebras can be used to unify diverse programming models through an axiomatization of commonalities, allowing proofs of fundamental program properties to be reused across program semantics that instantiate the algebra. The simple algebraic characterisation also facilities a straight-forward mechanisation of those proofs.

1.2. Wide-spectrum language

We motivate our semantic model and choice of primitive constructs by considering a range of specification and programming constructs they should be rich enough to express.

Preconditions. A specification c may be weakened by an assumption that its initial state satisfies a precondition p. In the sequential refinement calculus this is handled by a precondition command $\{p\}$ that terminates successfully if p holds, otherwise it aborts, i.e. it allows any behaviour whatsoever. The weakened specification is written "$\{p\} \,;\, c$", the sequential composition of $\{p\}$ and c. The key point is that an implementation (refinement) d of $\{p\} \,;\, c$ must behave as c when p holds initially but is unconstrained otherwise.

Postconditions. In the sequential refinement calculus a specification $(\!| q |\!)$, in which q is a binary relation, represents a requirement that if started in a state σ it must terminate in a state σ' such that $(\sigma, \sigma') \in q$. A similar construct is required in the context of concurrency but implementing such a specification is challenging because of possible environment interference (see Sect. 5.4). For this, stronger assumptions about the environment, such as rely conditions, are required (see below). The sequential specification $(\!| q |\!)$ only needs to consider the initial and final states and hence q is a binary relation. More generally, the execution of a program makes a sequence of steps and hence a specification could be generalised to constrain the sequence of steps. One such generalisation is to handle Jones' guarantee conditions (see below).

Rely conditions. To provide a compositional approach to refining concurrent programs, Jones [22, 23] introduced a rely condition, a binary relation r, that represents an assumption that the interference the environment makes between any two atomic steps of the program satisfies r between its before state and after state. Because there may be no environment steps between two program steps, Jones requires r to be reflexive, i.e. for all states σ, $(\sigma, \sigma) \in r$, and because there may be multiple environment steps between two program steps, Jones requires r to be transitive, i.e. $(\sigma, \sigma') \in r$ and $(\sigma', \sigma'') \in r$ imply $(\sigma, \sigma'') \in r$ (see Figure 1). Therefore, relation r is equal to its reflexive, transitive closure $r^\star$.

Handling rely conditions in the semantics requires the ability to refer explicitly to environment steps as well as program steps. The model presented in Section 2 is rich enough to allow one to define a command $(\epsilon\text{-}\mathbf{rely}\, r)$, that

[1] In other words, the operators of the calculus are *monotonic* in their arguments with respect to the refinement ordering.

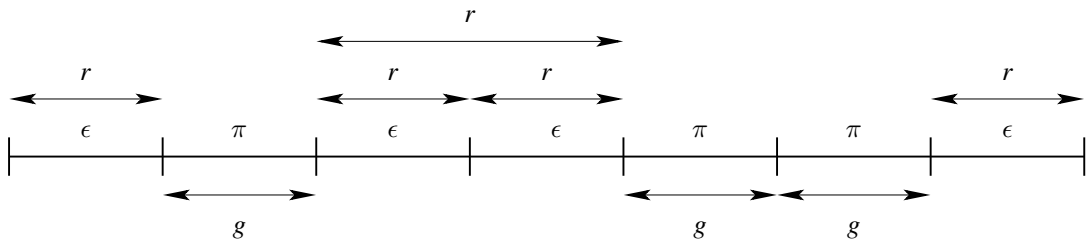

Fig. 1. Trace of execution with environment steps satisfying r and program steps satisfying g

represents an assumption that all environment steps satisfy the relation r. In the same way that a precondition command $\{p\}$ aborts if p does not hold initially, (ϵ-**rely** r) aborts if the environment makes a step not satisfying r (see Sect. 5.5). An environment assumption (ϵ-**rely** r) is added to a specification using the weak conjunction operator (see Sect. 4.5).

Guarantee conditions. If a program c_2 with a rely condition r is part of the environment of a program c_1, then c_1 should respect the rely condition r of c_2, i.e. every atomic program step taken by c_1 should satisfy r. To handle this, Jones introduced a guarantee condition — a binary relation g — and requires every atomic program step made by c_1 to satisfy g. As long as the guarantee g of c_1 is contained in the rely r of c_2, c_1 forms a suitable (part of the) environment for c_2.

A guarantee condition forms a restriction on the allowable steps made by a program. The semantic model is rich enough that one can define a guarantee command (π-**restrict** g), that forces every program step to satisfy the relation g (see Sect. 5.5). Such a guarantee can be combined with a specification using weak conjunction (see Sect. 4.5). Guarantee conditions may be generalised to allow more flexible constraints on the allowable sequences of steps.

Atomic specifications and linearisability. A *linearisable* operation [19] must (appear to) take place atomically between its invocation and response. It does not alter shared data outside of that single atomic step, although it may modify local data and perform other local computation (see Sect. 5.7). The framework straightforwardly handles atomic specifications, and distinguishes them from multi-step specifications such as postconditions described above.

Temporal logic. The semantic model is rich enough to encode temporal logics, such as linear temporal logic [27]. Because the semantics distinguishes program and environment steps, one can extend the temporal logic to distinguish between properties of the environment and of the program, as well as relating the two.

Expression evaluation. In the context of concurrency, the evaluation of an expression is subject to interference from concurrent programs modifying variables used within the expression and hence expression evaluation cannot be assumed to be atomic. Expressions are used in assignment statements and control structures such as conditionals (**if**) and repetitions (**while**). In the sequential refinement calculus a conditional control structure can be defined in terms of a *test* primitive $[b]$, that succeeds if b evaluates to true but is infeasible otherwise. It can be defined as follows, where "$\sqcap$" represents non-deterministic choice.

$$\textbf{if } b \textbf{ then } c_1 \textbf{ else } c_2 \mathrel{\widehat{=}} ([b] \,;\, c_1) \sqcap ([\neg b] \,;\, c_2)$$

Our approach preserves this form of definition by redefining the meaning of a test to be non-atomic and hence allow for interference during expression evaluation (see Sect. 5.1).

1.3. Contributions

The main contributions of this paper are (i) a semantics that supports concurrency for both programming language and specification constructs and hence reasoning about programs in a refinement calculus style, and (ii) a theory of atomic program and environment steps that treats environment steps as first-class citizens. The theory is based on a small set of primitive commands and operators, allowing the development of algebraic properties of operators. The semantics also forms an instance of a concurrent program algebra [13], which facilitates reasoning about properties of the model at a high level of abstraction. The theory supports rely and guarantee constructs, and it handles progress properties, e.g. temporal logic, including extensions to handle properties of the environment as well as the program.

Our semantic model, primitive commands, and program combinators are defined in Sections 2, 3 and 4, respectively. The wide-spectrum language is introduced in Section 5, and we discuss how it can be used to reason about progress properties in Section 6. We conclude in Section 7.

2. Semantic model

In our semantic model, program behaviours are described using traces of primitive steps. Section 2.1 introduces the states and the primitive steps used to construct these traces, and Section 2.2 defines the set of traces that are used to represent commands, along with some healthiness properties on sets of traces that represent commands.

2.1. States and primitive steps

To specify the behaviour of an executable command, it is sufficient to specify the allowable sequences of atomic steps (σ, σ') from state $\sigma \in \Sigma$ to state $\sigma' \in \Sigma$ made by the program, where Σ is the set of all possible program states. The semantics presented below is largely independent of the form of the set of program states Σ. One possible representation is as a store, i.e. a mapping from variables to values (including undefined), $\Sigma \mathrel{\widehat{=}} \mathit{Var} \to \mathit{Val}_\perp$, which we use here for illustration. Other representations are possible, for example, including a heap as well as a store.

Reactive-sequence semantics [7, 8] uses traces of the form $[(\sigma_1, \sigma_2), (\sigma_3, \sigma_4)]$, where the gap between σ_2 and σ_3 implicitly represents interference from the environment. In order to handle rely conditions, Peter Aczel [1, 6, 7, 8] needed to explicitly represent both environment and program steps, differentiating an environment step from σ to σ', written $\sigma \mathcal{E} \sigma'$ here, from a program step, written $\sigma \Pi \sigma'$. When steps are combined to form a trace, the post state σ_1 of each step $\sigma_0 \Pi \sigma_1$ must equal the pre state of the next step $\sigma_1 \Pi \sigma_2$ – and such a trace is said to be *consistent*. Because of this consistency, the two steps in sequence can be abbreviated to $\sigma_0 \Pi \sigma_1 \Pi \sigma_2$. For example, the following is a valid *Aczel trace*.

$$\sigma_0 \mathcal{E} \sigma_1 \Pi \sigma_2 \mathcal{E} \sigma_3 \Pi \sigma_4 \Pi \sigma_5 \mathcal{E} \sigma_6 \tag{1}$$

Note that environment steps are allowed before any program steps, $\sigma_0 \mathcal{E} \sigma_1$, and after all program steps, $\sigma_5 \mathcal{E} \sigma_6$. Aczel traces also allow multiple program steps without intervening environment steps and vice versa.[2]

Sequences of steps may be finite or infinite, and we use infinite Aczel traces to represent non-terminating behaviours. Finite traces represent behaviours that have either terminated successfully, aborted due to a failure of the program, aborted due to a failure of the environment, or become infeasible. These four cases are distinguished by including one of four special kinds of atomic step at the end of each finite trace. These special steps may only appear as the last step in a trace.

- Like De Roever [8], we use steps of the form $\sigma\checkmark$ to represent termination from some initial state σ. For example the Aczel trace $\sigma_0 \mathcal{E} \sigma_1 \Pi \sigma_2 \checkmark$, performs two steps and then terminates from the last state σ_2.
- To allow for aborting behaviour of programs, such as a precondition or rely condition failing to hold, the state space Σ is extended with a distinguished undefined state $\perp_s$, giving us $\Sigma_\perp \mathrel{\widehat{=}} \Sigma \mid \perp_s$. We then use step $\sigma \Pi \perp_s$ to represent the program taking an aborting step, and $\sigma \mathcal{E} \perp_s$ to represent the environment taking an aborting step, e.g. $\sigma_0 \mathcal{E} \sigma_1 \Pi \sigma_2 \Pi \perp_s$ or $\sigma_0 \mathcal{E} \sigma_1 \Pi \sigma_2 \mathcal{E} \perp_s$.
- A step from a state that does not lead anywhere, such as σ on its own, represents a step that is infeasible, e.g. the trace $\sigma_0 \mathcal{E} \sigma_1 \Pi \sigma_2$ performs two steps and then becomes infeasible from the last state σ_2. Programs may become infeasible right away, before taking any steps, e.g. the trace σ_0 is infeasible from its initial state σ_0. At first glance, the inclusion of infeasible traces may seem unnecessary, since they are, after all, unimplementable. However, in the same way that partial functions may be combined to create total functions, infeasible program behaviours may be combined to specify implementable ones, and the semantics benefits from the additional expressivity. For example, infeasibility may be used to represent the fact that a specification has conflicting requirements that cannot be met.

[2] De Roever [8] makes use of labels unique to each location and each thread in a program, rather than program (π) and environment (ϵ) markers. Here we follow Aczel's original simpler formulation more closely.

2.2. Traces and commands

We now formalise the Aczel traces that were introduced in the previous section. Fundamental definitions for traces used throughout the paper are contained in Figure 2. A trace – see (3) in Figure 2 – is denoted by an initial state $\sigma \in \Sigma$ and a possibly infinite sequence of atomic steps $t \in \text{seq}\, \mathit{Step}$, where we define

$$\mathit{Step} \mathrel{\widehat{=}} \Pi(\Sigma_\perp) \mid \mathcal{E}(\Sigma_\perp) \mid \checkmark \ .$$

For example, the trace written informally as $(\sigma_0\ \mathcal{E}\ \sigma_1\ \Pi\ \sigma_2\ \mathcal{E}\ \sigma_3\ \Pi\ \sigma_4\ \mathcal{E}\ \sigma_5\ \checkmark)$ is formalised as the trace

$$(\sigma_0, [\mathcal{E}(\sigma_1), \Pi(\sigma_2), \mathcal{E}(\sigma_3), \Pi(\sigma_4), \mathcal{E}(\sigma_5), \checkmark]) \ .$$

This form relies on Aczel traces being consistent. The *Step* definition does not need to include explicitly infeasible steps, since these are implicitly represented. For example, the trivial trace, $(\sigma, [\,])$, represents infeasible behaviour starting from σ.[3] A healthiness condition on traces is that a termination step, $\checkmark$, or an aborting step, either $\Pi(\perp_s)$ or $\mathcal{E}(\perp_s)$, can occur only as the last step of a (finite) trace (4).

We define a partial ordering, $\leq$, on traces that describes the conditions under which one trace $(\sigma_2, [t_2])$ is a valid refinement of another $(\sigma_1, [t_1])$ as follows.

$$(\sigma_1, [t_1]) \leq (\sigma_2, [t_2]) \mathrel{\widehat{=}} \sigma_1 = \sigma_2 \wedge \\ (t_2 \in \mathit{prefix}(t_1) \vee (\exists\, t_3 \in \mathit{Tseq} \cdot \mathit{finite}(t_3) \wedge t_1 = t_3 \frown \Pi(\perp_s) \wedge t_3 \in \mathit{prefix}(t_2)) \qquad (2)$$

Informally, this captures the notion that a trace may be refined by itself (since for any $t \in \mathit{Tseq}$, $t \in \mathit{prefix}(t)$), or by a trace that replaces a program abort step by any subsequent behaviour, or by a (prefix) trace that becomes infeasible. For example we have that

$$\begin{array}{ll} & (\sigma_0, [\Pi(\perp_s)]) \\ \leq & (\sigma_0, [\mathcal{E}(\sigma_1), \Pi(\perp_s)]) \\ \leq & (\sigma_0, [\mathcal{E}(\sigma_1), \Pi(\sigma_2), \Pi(\sigma_3), \mathcal{E}(\sigma_4), \checkmark]) \\ \leq & (\sigma_0, [\mathcal{E}(\sigma_1), \Pi(\sigma_2)]) \\ \leq & (\sigma_0, [\,]) \ . \end{array}$$

For each initial state σ_0, the set of traces that start from σ_0 form a complete lattice under the trace refinement ordering, with least element $(\sigma_0, [\Pi(\perp_s)])$, and greatest element $(\sigma_0, [\,])$. In this ordering there is an asymmetry to the interpretation of program abort steps and environment abort steps. The appearance of a $\Pi\perp_s$ step in a trace signifies that any behaviour whatsoever is acceptable from that step onwards (hence it may be arbitrarily refined), whereas $\mathcal{E}\perp_s$ does not: the only step tolerable at that point is environmental failure. Intuitively, this reflects our desire to distinguish an aborting program – one that may exhibit any possible behaviour – from one that will not fail unless that failure is brought about by its environment.

In our trace lattice for an initial σ_0, we have, for example, that

$$(\sigma_0, [\mathcal{E}(\sigma_1), \Pi(\sigma_2), \Pi(\sigma_3), \mathcal{E}(\sigma_4), \checkmark]) \sqcup (\sigma_0, [\mathcal{E}(\sigma_1), \mathcal{E}(\sigma_2), \Pi(\sigma_3), \checkmark]) \ = \ (\sigma_0, [\mathcal{E}(\sigma_1)])$$

where $\sqcup$ is the supremum operator. Note how the trace becomes infeasible at the first incompatible step.

The trace ordering is lifted pointwise on initial states to define an ordering on the set of commands themselves. The set of commands, (10) in Figure 2, are denoted by sets of traces with additional healthiness requirements imposed by the *prefix_closure*(..) and *abort_closure*(..) conditions. These conditions are equivalent to requiring that each command has at least one trace $(\sigma, [\,])$, starting from each possible initial state σ, and that they are up-closed with respect to the trace ordering (2). This allows us to define refinement of commands by trace inclusion as per usual: for commands c and d we have that

$$c \sqsubseteq d \mathrel{\widehat{=}} c \supseteq d \ .$$

In more detail, the empty closure (5) of a set of traces, s, adds all traces of the form $(\sigma, [\,])$ for all $\sigma \in \Sigma$ and the prefix closure (6) adds all prefixes of the traces of s as well. The aborting traces of s are those ending in $\Pi(\perp_s)$ but with that last step removed (7). Abort completing a set of (aborting) traces adds all possible extensions of the traces (8) and the abort closure of a set of traces completes all its aborting traces (9). Note that the abort closure only applies

[3] Here we differ from both Aczel [1] and De Roever [8], who do not allow empty traces.

Let s be a set of traces. For a sequence t, $\mathrm{dom}(t)$ stands for the set of indices of t, which may be infinite if t is an infinite sequence.

$$Tr \mathrel{\widehat=} \Sigma \times Tseq \tag{3}$$
$$Tseq \mathrel{\widehat=} \{t \in \mathrm{seq}\, Step \mid \forall i \in \mathrm{dom}(t) \cdot t(i) \in \{\Pi(\bot_s), \mathcal{E}(\bot_s), \checkmark\} \Rightarrow i+1 \notin \mathrm{dom}(t))\} \tag{4}$$
$$empty_closure(s) \mathrel{\widehat=} s \cup \{(\sigma, t) \in Tr \mid t = [\,]\} \tag{5}$$
$$prefix_closure(s) \mathrel{\widehat=} empty_closure(s) \cup \{(\sigma, t) \in Tr \mid (\exists\, t' \cdot (\sigma, t') \in s \wedge t \in prefix(t'))\} \tag{6}$$
$$aborting(s) \mathrel{\widehat=} \{(\sigma, t) \in Tr \mid finite(t) \wedge (\sigma, t \frown [\Pi(\bot_s)]) \in s\} \tag{7}$$
$$abort_complete(s) \mathrel{\widehat=} \{(\sigma, t) \in Tr \mid (\exists\, t' \in prefix(t) \cdot (\sigma, t') \in s)\} \tag{8}$$
$$abort_closure(s) \mathrel{\widehat=} s \cup abort_complete(aborting(s)) \tag{9}$$
$$Com \mathrel{\widehat=} \{s \in \mathbb{P}\, Tr \mid s = prefix_closure(s) \wedge s = abort_closure(s)\} \tag{10}$$

Fig. 2. Traces and associated auxiliary functions

Let r be a binary relation on states, $r \subseteq \Sigma \times \Sigma$, and p be a set of states, $p \subseteq \Sigma$.

$$\boldsymbol{\pi}(r) \mathrel{\widehat=} prefix_closure(\{(\sigma, \sigma') \in r \cdot (\sigma, [\Pi(\sigma'), \checkmark])\}) \tag{11}$$
$$\boldsymbol{\epsilon}(r) \mathrel{\widehat=} prefix_closure(\{(\sigma, \sigma') \in r \cdot (\sigma, [\mathcal{E}(\sigma'), \checkmark])\}) \tag{12}$$
$$\boldsymbol{\tau}(p) \mathrel{\widehat=} empty_closure(\{\sigma \in p \cdot (\sigma, [\checkmark])\}) \tag{13}$$
$$\bot_{\boldsymbol{\pi}} \mathrel{\widehat=} Tr \tag{14}$$
$$\bot_{\boldsymbol{\epsilon}} \mathrel{\widehat=} empty_closure(\{\sigma \in \Sigma \cdot (\sigma, [\mathcal{E}(\bot_s)])\}) \tag{15}$$
$$\boldsymbol{\epsilon}_\bot(r) \mathrel{\widehat=} \boldsymbol{\epsilon}(r) \cup \bot_{\boldsymbol{\epsilon}} \tag{16}$$

Fig. 3. Semantics of primitive commands

to traces ending in $\Pi(\bot_s)$ and not traces ending in $\mathcal{E}(\bot_s)$ for the reasons discussed above. Commands are represented by sets of traces that are both prefix closed (including empty closed) and abort closed (10).[4]

3. Primitive commands

A novelty of our approach is that the semantics of commands is defined in terms of a small set of primitive commands plus a set combinators, such as sequential and parallel composition. The formal definition of the primitive commands is given in Figure 3 in which the notation $\{x \in S \mid p(x) \cdot e(x)\}$ stands for the set of all the values of the expression $e(x)$ for the bound variable x ranging over the values in the set S such that the predicate $p(x)$ holds, $\{x \in S \cdot e(x)\}$ abbreviates $\{x \in S \mid \mathsf{true} \cdot e(x)\}$, and $\{x \in S \mid p(x)\}$ abbreviates $\{x \in S \mid p(x) \cdot x\}$. Aczel's program and environment steps motivate the first two primitive commands, and De Roever's terminating step motivates the third.

- $\boldsymbol{\pi}(r)$ can make a single program step $\sigma \Pi \sigma'$ satisfying binary relation r, i.e. $(\sigma, \sigma') \in r$, and terminate (11).
- $\boldsymbol{\epsilon}(r)$ can make a single environment step $\sigma \mathcal{E} \sigma'$ satisfying r and terminate (12).
- $\boldsymbol{\tau}(p)$ can make a terminating step $\sigma\checkmark$ from any state σ in the set of states p (13).

The prefix and empty closures in definitions (11)–(13) ensure the commands satisfy the healthiness condition of being abort and prefix (and hence empty) closed. Our final three primitive commands handle aborting behaviours.

- $\bot_{\boldsymbol{\pi}}$ allows any behaviour at all (14). It equals the abort and prefix-closure of the set $\{\sigma \in \Sigma \cdot (\sigma, [\Pi(\bot_s)])\}$.
- $\bot_{\boldsymbol{\epsilon}}$ allows the environment to abort (15).
- $\boldsymbol{\epsilon}_\bot(r)$ allows the environment to either abort or do a step $\sigma \mathcal{E} \sigma'$ satisfying r and terminate (16).

[4] From the perspective of mechanisation using a theorem prover, we find the definitions of prefix closure and abort closure more practical than up-closing directly with respect to the trace ordering (2).

We introduce the following abbreviations for the common cases where the (program or environment step) relation is the universal relation, $\mathsf{univ} \mathrel{\widehat{=}} \Sigma \times \Sigma$, that relates all pairs of states, and where the test set is either the set of all states, Σ, or the empty set, $\varnothing$:

$$\begin{array}{lcllcllcl} \boldsymbol{\pi} & \mathrel{\widehat{=}} & \boldsymbol{\pi}(\mathsf{univ}) & \quad \boldsymbol{\epsilon} & \mathrel{\widehat{=}} & \boldsymbol{\epsilon}(\mathsf{univ}) & \quad \boldsymbol{\epsilon}_\bot & \mathrel{\widehat{=}} & \boldsymbol{\epsilon}_\bot(\mathsf{univ}) \\ \boldsymbol{\tau} & \mathrel{\widehat{=}} & \boldsymbol{\tau}(\Sigma) & \quad \top & \mathrel{\widehat{=}} & \boldsymbol{\tau}(\varnothing) \end{array}$$

The command $\boldsymbol{\pi}$ can make any single program step and terminate; $\boldsymbol{\epsilon}$ can make any single environment step and terminate, and $\boldsymbol{\epsilon}_\bot$ allows the environment to abort or take any environment step and terminate. Command $\boldsymbol{\tau}$ terminates immediately from any initial state, while $\top$ is infeasible from all initial states, having no traces other than $(\sigma, [\,])$ for all states $\sigma \in \Sigma$. Note that $\boldsymbol{\pi}(\varnothing) = \boldsymbol{\epsilon}(\varnothing) = \top$.

If $r_2 \subseteq r_1$, the refinements $\boldsymbol{\pi}(r_1) \sqsubseteq \boldsymbol{\pi}(r_2)$ and $\boldsymbol{\epsilon}(r_1) \sqsubseteq \boldsymbol{\epsilon}(r_2)$ hold, and if $p_2 \subseteq p_1$, $\boldsymbol{\tau}(p_1) \sqsubseteq \boldsymbol{\tau}(p_2)$. Commands form a complete lattice under the refinement ordering with $\bot_{\boldsymbol{\pi}}$ as the least element and $\top$ as the greatest element, that is, for any c, $\bot_{\boldsymbol{\pi}} \sqsubseteq c \sqsubseteq \top$.

4. Program combinators

This section defines the primitive program combinators found in Figure 4 (using the auxiliary functions defined in both Figures 2 and 4). The choice operators (the infimum and supremum of the program lattice) can be found in Section 4.1, sequential composition in Section 4.2, and Section 4.3 uses these to define iteration constructs using fixed points. Section 4.4 defines parallel composition and Section 4.5 a weak conjunction operator [24] that is used in Section 5.5 to define constructs suitable for specifying relies and guarantees. Section 4.6 introduces a program variable unrestriction operator used in the definition of local variable blocks.

4.1. Infimum and supremum

A nondeterministic choice (or *infimum*), $\bigsqcap C$, over a set C of commands gives the union of the behaviour of its components (17). In this definition, the empty closure ensures that the choice is a valid command, even if the choice is over the empty set. Supremum ($\bigsqcup$) in the lattice of commands is represented by intersection of sets of traces.

As binary operators $\sqcap$ (18) and $\sqcup$ (20) are associative, commutative and idempotent. Operator $\sqcap$ has the identity $\top$ (i.e. $\top \sqcap c = c$) and annihilator $\bot_{\boldsymbol{\pi}}$ (i.e. $\bot_{\boldsymbol{\pi}} \sqcap c = \bot_{\boldsymbol{\pi}}$), and $\sqcup$ has the identity $\bot_{\boldsymbol{\pi}}$ and annihilator $\top$. When applied to primitive commands, the choice operators satisfy the following properties.

$$\begin{array}{rcllcl} \boldsymbol{\tau}(p_1) \sqcap \boldsymbol{\tau}(p_2) & = & \boldsymbol{\tau}(p_1 \cup p_2) & \text{and} & \boldsymbol{\tau}(p_1) \sqcup \boldsymbol{\tau}(p_2) & = & \boldsymbol{\tau}(p_1 \cap p_2) \\ \boldsymbol{\pi}(r_1) \sqcap \boldsymbol{\pi}(r_2) & = & \boldsymbol{\pi}(r_1 \cup r_2) & \text{and} & \boldsymbol{\pi}(r_1) \sqcup \boldsymbol{\pi}(r_2) & = & \boldsymbol{\pi}(r_1 \cap r_2) \\ \boldsymbol{\epsilon}_\bot(r_1) \sqcap \boldsymbol{\epsilon}_\bot(r_2) & = & \boldsymbol{\epsilon}_\bot(r_1 \cup r_2) & \text{and} & \boldsymbol{\epsilon}_\bot(r_1) \sqcup \boldsymbol{\epsilon}_\bot(r_2) & = & \boldsymbol{\epsilon}_\bot(r_1 \cap r_2) \end{array}$$

4.2. Sequential composition

Sequential composition $c_1 \,;\, c_2$, written $c_1 c_2$ for short, retains all the unterminated traces of c_1 (these include infeasible/incomplete (prefix) traces, aborting traces and infinite traces) as well as the terminated traces $tr_1 \in c_1$ concatenated with traces $tr_2 \in c_2$ that start from the last state of tr_1 (21). The definition makes use of the set of terminating traces of a command (28); note that the final "✓" step is removed from the terminating traces as part of this definition. For two sets of traces s_1 and s_2, for which the traces of s_1 are finite, $s_1 \frown s_2$ forms the set of concatenations of traces $tr_1 \in s_1$ and $tr_2 \in s_2$ such that the last state (31) of tr_1 matches the initial state of tr_2 (29). The abort closure of the unterminated traces of c_1 is used to ensure the set of traces of the sequential composition is abort closed.

Sequential composition is associative. It has identity $\boldsymbol{\tau}$ and left annihilators of $\top$ and $\bot_{\boldsymbol{\pi}}$. However note that in general, $(c \,\bot_{\boldsymbol{\pi}}) \neq \bot_{\boldsymbol{\pi}}$ and $(c \,\top) \neq \top$, for example, if c has no terminating traces. Sequential composition distributes over nondeterministic choice from the right: $(\bigsqcap C) \,;\, d = \bigsqcap \{c \in C \cdot c \,;\, d\}$ but only over finite choices from the left: $c \,;\, (d_1 \sqcap d_2) = (c \,;\, d_0) \sqcap (c \,;\, d_2)$. Sequential compositions of primitive commands satisfy the following properties, where $p \lhd r$ is the relation r restricted so that its domain is contained in p.

$$\boldsymbol{\tau}(p_1)\, \boldsymbol{\tau}(p_2) = \boldsymbol{\tau}(p_1 \cap p_2)$$

Let C be a set of commands, c, c_1 and c_2 be commands and x a variable name.

$$\bigsqcap C \mathrel{\widehat=} \mathit{empty_closure}(\bigcup C) \tag{17}$$

$$c_1 \sqcap c_2 \mathrel{\widehat=} \bigsqcap\{c_1, c_2\} \;=\; c_1 \cup c_2 \tag{18}$$

$$\bigsqcup C \mathrel{\widehat=} (\bigcap C) \tag{19}$$

$$c_1 \sqcup c_2 \mathrel{\widehat=} \bigsqcup\{c_1, c_2\} \;=\; c_1 \cap c_2 \tag{20}$$

$$c_1 \,;\, c_2 \mathrel{\widehat=} \mathit{abort_closure}(c_1 - \mathit{terminating}(c_1)) \cup (\mathit{terminating}(c_1) \frown c_2) \tag{21}$$

$$c^\star \mathrel{\widehat=} (\nu x \cdot \boldsymbol{\tau} \sqcap c\, x) \tag{22}$$

$$c^\omega \mathrel{\widehat=} (\mu x \cdot \boldsymbol{\tau} \sqcap c\, x) \tag{23}$$

$$c^\infty \mathrel{\widehat=} c^\omega \top \tag{24}$$

$$c_1 \parallel c_2 \mathrel{\widehat=} \mathit{abort_closure}(\{tr \in \mathit{Tr} \mid (\exists\, tr_1 \in c_1, tr_2 \in c_2 \cdot \mathit{match_trace}(tr_1, tr_2, tr))\}) \tag{25}$$

$$c_1 \Cap c_2 \mathrel{\widehat=} (c_1 \cap c_2) \cup \mathit{abort_first}(c_1, c_2) \cup \mathit{abort_first}(c_2, c_1) \tag{26}$$

$$c \setminus x \mathrel{\widehat=} \{(\sigma, t) \in \mathit{Tr} \mid \exists(\sigma', t') \in c \cdot x \lhd (\sigma, t) = x \lhd (\sigma', t')\} \tag{27}$$

where for s, s_1 and s_2 sets of traces, z_1 and z_2 steps, σ, σ_1 and σ_2 states, and t, t_1 and t_2 traces,

$$\mathit{terminating}(s) \mathrel{\widehat=} \{(\sigma, t) \in \mathit{Tr} \mid \mathit{finite}(t) \wedge (\sigma, t \frown [\checkmark]) \in s\} \tag{28}$$

$$s_1 \frown s_2 \mathrel{\widehat=} \{\; (\sigma_1, t_1) \in s_1;\; (\sigma_2, t_2) \in s_2 \mid \mathit{finite}(t_1) \wedge \mathit{last_state}(\sigma_1, t_1) = \sigma_2 \cdot (\sigma, t_1 \frown t_2)\} \tag{29}$$

$$\mathit{last_state}(\sigma, [\,]) = \sigma \tag{30}$$

$$\mathit{last_state}(\sigma, t \frown [z]) = \sigma' \quad \text{if } z \in \{\Pi(\sigma'), \mathcal{E}(\sigma')\} \tag{31}$$

$$\mathit{match_step}(z_1, z_2, z) \mathrel{\widehat=} \exists\, \sigma \in \Sigma_\perp \cdot \; \begin{array}{l} z_1 = \Pi(\sigma) \wedge z_2 = \mathcal{E}(\sigma) \wedge z = \Pi(\sigma) \vee \\ z_1 = \mathcal{E}(\sigma) \wedge z_2 = \Pi(\sigma) \wedge z = \Pi(\sigma) \vee \\ z_1 = \mathcal{E}(\sigma) \wedge z_2 = \mathcal{E}(\sigma) \wedge z = \mathcal{E}(\sigma) \vee \\ z_1 = \checkmark \wedge z_2 = \checkmark \wedge z = \checkmark \end{array} \tag{32}$$

$$\mathit{match_trace}((\sigma_1, t_1), (\sigma_2, t_2), (\sigma, t)) \mathrel{\widehat=} \sigma_1 = \sigma_2 = \sigma \wedge \mathrm{dom}(t_1) = \mathrm{dom}(t_2) = \mathrm{dom}(t) \wedge \tag{33}$$

$$(\forall\, i \in \mathrm{dom}(t) \cdot \mathit{match_step}(t_1(i), t_2(i), t(i))) \tag{34}$$

$$\mathit{abort_first}(c_1, c_2) \mathrel{\widehat=} \mathit{abort_complete}(\mathit{aborting}(c_1) \cap c_2) \tag{35}$$

$$x \lhd (\sigma, t) \mathrel{\widehat=} (\{x\} \lhd \sigma, (\lambda\, i \in \mathrm{dom}(t) \cdot \{x\} \lhd t(i))) \tag{36}$$

Fig. 4. Semantics of primitive operators

$$\boldsymbol{\tau}(p)\, \boldsymbol{\pi}(r) = \boldsymbol{\pi}(p \lhd r)$$
$$\boldsymbol{\tau}(p)\, \boldsymbol{\epsilon}_\perp(r) = \boldsymbol{\epsilon}_\perp(p \lhd r)$$

Notationally we give sequential composition a higher precedence than all the other binary operators.

Using nondeterministic choice and sequential composition one can define the precondition command $\{p\}$ (37); a command, $\mathit{opt}(r)$, that allows a program step $\boldsymbol{\pi}(r)$ to be optional for states σ such that $(\sigma, \sigma) \in r$ by using a silent τ step instead in this case, thus allowing code optimisations that remove redundant steps (38); and an atomic update of a variable x to a constant κ, that is defined using opt so that it can be optimised to the null command if x already equals κ (39). The notation $\sigma[x \backslash \kappa]$ stands for the state σ with the value of x updated to be κ.

$$\{p\} \mathrel{\widehat=} \boldsymbol{\tau}(p) \sqcap (\boldsymbol{\tau}(\overline{p})\, \perp_{\boldsymbol{\pi}}) \tag{37}$$

$$\mathit{opt}(r) \mathrel{\widehat=} \boldsymbol{\pi}(r) \sqcap \boldsymbol{\tau}(\{\sigma \in \Sigma \mid (\sigma, \sigma) \in r\}) \tag{38}$$

$$\mathit{update}(x, \kappa) \mathrel{\widehat=} \mathit{opt}(\{(\sigma, \sigma') \in \Sigma \times \Sigma \mid \sigma' = \sigma[x \backslash \kappa]\}) \tag{39}$$

The precondition command $\{p\}$ terminates immediately from states in p but aborts from states in the complement $\overline{p}$ of p. Note that $\{\Sigma\} = \boldsymbol{\tau}(\Sigma) \sqcap (\boldsymbol{\tau}(\varnothing)\, \perp_{\boldsymbol{\pi}}) = \boldsymbol{\tau} \sqcap (\top\, \perp_{\boldsymbol{\pi}}) = \boldsymbol{\tau}$ and $\{\varnothing\} = \boldsymbol{\tau}(\varnothing) \sqcap (\boldsymbol{\tau}(\Sigma)\, \perp_{\boldsymbol{\pi}}) = \top \sqcap (\boldsymbol{\tau}\, \perp_{\boldsymbol{\pi}}) = \perp_{\boldsymbol{\pi}}$.

4.3. Recursion and iteration

The set of all commands under the refinement ordering forms a complete lattice and hence least and greatest fixed points can be defined over monotone functions from commands to commands. For a monotonic function f from commands to commands, μf is its least fixed point and νf is its greatest fixed point. As usual $\mu(\lambda x \cdot f(x))$ is abbreviated as $\mu x \cdot f(x)$, and likewise for greatest fixed points.

The iteration operators are defined in terms of least (μ) and greatest (ν) fixed points with respect to the refinement ordering. The command $c^\star$ iterates c a finite number of times, zero or more (22); c^ω iterates c any number of times, zero or more and possibly infinitely many times (23); and c^∞ iterates c an infinite number of times (24). The only feasible traces of c^∞ are the infinite traces of c^ω.

A number of distinguished commands such as **skip** and **idle** can be defined using iteration and the programming primitives. Whereas $\boldsymbol{\tau}$ terminates immediately with no environment or program steps, the command **skip** allows any environment steps but no program steps, and the command **idle** allows any environment steps as well as a finite number of stuttering program steps (i.e. steps that do not change the state); it represents the no-op command of code.[5] The relation id is the identity relation on states.

$$\textbf{skip} \;\widehat{=}\; \boldsymbol{\epsilon}^\omega_\bot \qquad\qquad \textbf{idle} \;\widehat{=}\; (\boldsymbol{\pi}(\mathsf{id}) \sqcap \boldsymbol{\epsilon}_\bot)^\star\, \boldsymbol{\epsilon}^\omega_\bot$$

Note that although $(\boldsymbol{\tau}\, c) = c$, it is not the case that $(\textbf{skip}\, c) = c$ in general because **skip** allows environment steps that may change the state. The command **idle** is refined by **skip** (and is equivalent to **skip** modulo finite stuttering) and **skip** is refined by $\boldsymbol{\tau}$, i.e. $\textbf{idle} \sqsubseteq \textbf{skip} \sqsubseteq \boldsymbol{\tau}$.

Two further distinguished commands are **chaos** and **term**. The command **chaos** represents a program that can do any non-aborting behaviour and allows its environment to do anything, including abort. The command **term** represents terminating program behaviour (but can't force its environment to terminate), i.e. if $\textbf{term} \sqsubseteq c$ then c has finitely many program steps. Termination is discussed in more detail in Section 5.3.

$$\textbf{chaos} \;\widehat{=}\; (\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot)^\omega \qquad\qquad \textbf{term} \;\widehat{=}\; (\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot)^\star\, \boldsymbol{\epsilon}^\omega_\bot$$

A preempted computation is one ending in an infinite sequence of environment steps. Note that the command **term** includes preempted behaviour ($\textbf{term} \sqsubseteq \textbf{preempted}$) which reflects the notion that an implementation may not terminate if the environment takes over forever. The command **forever** represents all infinite behaviours. Any non-aborting behaviour is either terminating or infinite: $\textbf{chaos} = \textbf{term} \sqcap \textbf{forever}$. Such properties can be proved from the algebraic properties of iterations [29], such as $c^\omega = c^\star \sqcap c^\infty$ and $(c \sqcap d)^\omega = (d^\omega\, c)^\omega\, d^\omega$.

$$\textbf{preempted} \;\widehat{=}\; (\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot)^\star\, \boldsymbol{\epsilon}^\infty_\bot \qquad\qquad \textbf{forever} \;\widehat{=}\; (\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot)^\infty$$

4.4. Parallel composition

Parallel composition, $c_1 \parallel c_2$, synchronises a program step of c_1 with an environment step of c_2 to give a program step of $c_1 \parallel c_2$ (and vice versa), and synchronises an environment step of c_1 with an environment step of c_2 to give an environment step of $c_1 \parallel c_2$ – see *match_step* (32) in Figure 4. For example,

$$\begin{array}{ll} (\sigma_0\ \mathcal{E}\ \sigma_1\ \Pi\ \sigma_2\ \Pi\ \sigma_3\ \mathcal{E}\ \sigma_4\ \checkmark) & \text{of } c_1 \text{ in parallel with} \\ (\sigma_0\ \Pi\ \sigma_1\ \mathcal{E}\ \sigma_2\ \mathcal{E}\ \sigma_3\ \mathcal{E}\ \sigma_4\ \checkmark) & \text{of } c_2 \text{ gives} \\ (\sigma_0\ \Pi\ \sigma_1\ \Pi\ \sigma_2\ \Pi\ \sigma_3\ \mathcal{E}\ \sigma_4\ \checkmark) & \text{of } c_1 \parallel c_2. \end{array}$$

Note that program steps can only be synchronised with a corresponding environment step and cannot synchronise with another program step. The synchronisations allow steps with after state $\bot_s$ allowing synchronisation of aborting program and environment steps or aborting environment steps of both programs. For parallel composition a $\sigma\Pi\bot_s$ step is matched by the environment step $\sigma\mathcal{E}\bot_s$ to give the program step $\sigma\Pi\bot_s$ of the parallel composition. Because a $\sigma\Pi\bot_s$ step must be matched by a $\sigma\mathcal{E}\bot_s$ step, a (rather strong) form of specification is one that rules out $\sigma\mathcal{E}\bot_s$ steps. This gives a hint to the expressive power of the semantics.

Our approach differs from that of De Roever [8] in that we require both programs to terminate together, whereas he allows one program to terminate early and the parallel composition to become the second program, and hence he uses $\boldsymbol{\tau}$ as the identity of parallel composition. Our approach is required to handle parallel composition of end-to-end specifications (see Sect. 5.4). Our definition of parallel composition does not include any notion of fairness (see

[5] It is similar to the **idle** command in the real-time context that changes nothing but allows time to pass [15, 17].

Sect. 6). Parallel composition involving primitives satisfies simple rules such as the following.

$$(\boldsymbol{\tau}(p_1)\,c_1) \parallel (\boldsymbol{\tau}(p_2)\,c_2) = \boldsymbol{\tau}(p_1 \cap p_2)\,(c_1 \parallel c_2)$$
$$(\boldsymbol{\pi}(r_1)\,c_1) \parallel (\boldsymbol{\epsilon}_\bot(r_2)\,c_2) = \boldsymbol{\pi}(r_1 \cap r_2)\,(c_1 \parallel c_2)$$
$$(\boldsymbol{\epsilon}_\bot(r_1)\,c_1) \parallel (\boldsymbol{\epsilon}_\bot(r_2)\,c_2) = \boldsymbol{\epsilon}_\bot(r_1 \cap r_2)\,(c_1 \parallel c_2)$$

Parallel composition is associative and commutative. Because **skip** allows an arbitrary number of environment steps that may match any steps of c, $c \parallel \mathbf{skip} = c$, and hence **skip** is the identity of parallel composition. Note that **skip** allows an infinite sequence of environment steps; this allows it to synchronise with infinite traces of c. Note that because $\boldsymbol{\tau}$ does not allow environment steps, it is not the identity of parallel composition, in fact $c \parallel \boldsymbol{\tau}$ can terminate only if c behaves as $\boldsymbol{\tau}$.

A desirable property of parallel composition is that an infeasible program, e.g. $\top$, in parallel with any command c, including $\bot_{\boldsymbol{\pi}}$, is infeasible. In order to ensure $\top$ (which can make no steps at all) annihilates $\bot_{\boldsymbol{\pi}}$, parallel composition requires that an aborting program step $\sigma \Pi \bot_s$ is matched by an environment step $\sigma \mathcal{E} \bot_s$. Because $\top$ doesn't include the step $\sigma \mathcal{E} \bot_s$ (it has no steps at all), $\top \parallel \bot_{\boldsymbol{\pi}} = \top$. More generally, parallel composition distributes over nondeterministic choice: $(\bigsqcap C) \parallel d = \bigsqcap \{c \in C \cdot c \parallel d\}$.

4.5. Weak conjunction

Although the supremum operator allows us to conjoin commands, i.e. $c_1 \sqcup c_2$ forms the intersection of the traces of c_1 and c_2 and hence is constrained to satisfy the specifications of both c_1 and c_2, it is often too restrictive for our purposes. Most programs fulfill commitments under given assumptions, and can fail if those assumptions are not met. Often, one would like to constrain such a specification, c_1 say, so that it satisfies an additional commitment, c_2, only when c_1's assumptions are met (i.e. it does not abort). We use a weak conjunction operator for that purpose, and this allows us to handle rely and guarantee conditions is a general way.[6] The weak conjunction of two commands, $c_1 \Cap c_2$, synchronises non-aborting steps, e.g.,

$$\begin{array}{ll} (\sigma_0\ \mathcal{E}\ \sigma_1\ \Pi\ \sigma_2\ \Pi\ \sigma_3\ \checkmark) & \text{of } c_1 \text{ combines with} \\ (\sigma_0\ \mathcal{E}\ \sigma_1\ \Pi\ \sigma_2\ \Pi\ \sigma_3\ \checkmark) & \text{of } c_2 \text{ to give} \\ (\sigma_0\ \mathcal{E}\ \sigma_1\ \Pi\ \sigma_2\ \Pi\ \sigma_3\ \checkmark) & \text{of } c_1 \Cap c_2, \end{array}$$

but if either performs a program abort step, it aborts, for example,

$$\begin{array}{ll} (\sigma_0\ \Pi\ \sigma_1\ \mathcal{E}\ \sigma_2\ \Pi\ \bot_s) & \text{of } c_1 \text{ combines with} \\ (\sigma_0\ \Pi\ \sigma_1\ \mathcal{E}\ \sigma_2) & \text{of } c_2 \text{ to give} \\ (\sigma_0\ \Pi\ \sigma_1\ \mathcal{E}\ \sigma_2\ \Pi\ \bot_s) & \text{of } c_1 \Cap c_2. \end{array}$$

The traces of c_2 include infeasible traces (and their prefixes) so that in the weak conjunction $c_1 \Cap c_2$, a trace of c_1 that leads to an aborting step is matched by a trace of c_2, even if there are no feasible extensions of the trace of c_2. A weak conjunction $c_1 \Cap c_2$ contains the intersection of the traces of c_1 and c_2 as well as the aborting traces of one for which there is a trace of the other that matches up until (but not including) the final abort step (26). As the sets of traces are prefix closed, any prefixes of c_1 and c_2 which match are included. To handle weak conjunction, care must be taken with commands that include infeasible components. In the sequential refinement calculus for any terminating command c,

$$c\,\top = \top. \tag{40}$$

This property does not hold for the semantics given here because

$$(c\,\bot_{\boldsymbol{\pi}}) \Cap (c\,\top) \sqsubseteq c\,(\bot_{\boldsymbol{\pi}} \Cap \top) = c\,\bot_{\boldsymbol{\pi}} \tag{41}$$

Were we to allow (40) the left side reduces to

$$c\,\bot_{\boldsymbol{\pi}} \Cap \top$$

which assuming c is not $\bot_{\boldsymbol{\pi}}$ reduces to $\top$.

[6] Some earlier papers refer to "$\Cap$" as strict conjunction (because it is abort strict). This led to some confusion because it is not a strong as conjunction (intersection of traces).

Weak conjunction is associative, commutative and idempotent. It has identity **chaos** and $\bot_{\boldsymbol{\pi}}$ is an annihilator. For example, one can deduce the following.

$$\begin{array}{rcl}(\boldsymbol{\tau}(p_1)\,c_1) \mathbin{⫛} (\boldsymbol{\tau}(p_2)\,c_2) &=& \boldsymbol{\tau}(p_1 \cap p_2)\,(c_1 \mathbin{⫛} c_2)\\ (\boldsymbol{\pi}(r_1)\,c_1) \mathbin{⫛} (\boldsymbol{\pi}(r_2)\,c_2) &=& \boldsymbol{\pi}(r_1 \cap r_2)\,(c_1 \mathbin{⫛} c_2)\\ (\boldsymbol{\epsilon}_\bot(r_1)\,c_1) \mathbin{⫛} (\boldsymbol{\epsilon}_\bot(r_2)\,c_2) &=& \boldsymbol{\epsilon}_\bot(r_1 \cap r_2)\,(c_1 \mathbin{⫛} c_2)\end{array} \qquad \begin{array}{rcl} c \mathbin{⫛} \bot_{\boldsymbol{\pi}} &=& \bot_{\boldsymbol{\pi}}\\ (\boldsymbol{\pi}(r)\,c_0) \mathbin{⫛} (\boldsymbol{\epsilon}_\bot\,c_1) &=& \top\end{array}$$

4.6. Unrestricting a variable

In order to define local variable blocks (in Section 5.6) the operator $c \setminus x$ is used to define a program with traces that are the same as c, except they do not need to agree on values of the variable x. Foster et al. use a similar operator in their work [12, Sect. 5.4]. Let $x \lhd (\sigma, t)$ be the trace (σ, t) with variable x removed from σ and from the state of every program and environment step in t. The traces of the command $c \setminus x$ are the same as those of c, except they do not restrict the variable x at all (27).

5. A wide-spectrum language

In this section we extend the basic language to an imperative concurrent programming language augmented with specification constructs. Section 5.1 addresses the semantics of expression evaluation without the common assumption that expression evaluation is atomic. Section 5.2 make use of expressions to define typical imperative programming constructs: assignment, "if" statement and "while" statement. Section 5.3 addresses specifying termination and fairness constraints. End-to-end specifications are introduced in Section 5.4; they are the generalisation of postcondition specifications to the concurrent context. Section 5.5 introduces rely and guarantee constructs that allow the language to express Jones-style relational rely and guarantee specifications. Local variable blocks are illustrated in Section 5.6.

5.1. Expression evaluation

In the context of concurrency, expression evaluation is subject to interference on shared variables and hence cannot be treated as being atomic. As a consequence of this, expression evaluation is nondeterministic because the interference may modify variables during evaluation. Atomic read and write of shared variables is however assumed; (structured) variables that cannot be read/written atomically need to be treated separately.

The syntax of expressions is standard, encompassing values, variables, and the typical unary and binary operators (to save space we do not consider generalisation to n-ary operators). Although side effects within expressions can be handled by the semantics, only expressions free from side effects are considered here and hence any program steps made during expression evaluation are stuttering steps, i.e. they do not modify the observable state. To allow for any possible implementation strategy, the set of traces of an expression evaluation is closed under finite stuttering, i.e. if it includes a trace t_1, it includes any other trace t_2 that is equivalent to t_1 modulo finite stuttering of program steps. To allow for undefined expressions, the set Val is augmented with the undefined value $\bot_v$ to give the set $\mathit{Val}_\bot$.

The command $[\![e]\!]_\kappa$ represents the evaluation of the expression e to the value $\kappa \in \mathit{Val}_\bot$. Evaluation of e to κ may either succeed or fail. To model this in the semantics an evaluation that succeeds is terminating whereas a failing evaluation is infeasible and hence is eliminated from a nondeterministic choice with a succeeding evaluation to some other value. The semantics of expression evaluation is defined below, where x is a variable, $\ominus$ is a unary operator and $\oplus$ is a binary operator.

$$[\![\kappa_1]\!]_{\kappa_2} \mathrel{\widehat{=}} \mathbf{idle}\,\boldsymbol{\tau}(\{\sigma \in \Sigma \mid \kappa_1 = \kappa_2\}) \qquad \text{for } \kappa_1, \kappa_2 \in \mathit{Val}_\bot \tag{42}$$

$$[\![x]\!]_{\kappa} \mathrel{\widehat{=}} \mathbf{idle}\,\boldsymbol{\tau}(\{\sigma \in \Sigma \mid \sigma(x) = \kappa\})\,\mathbf{idle} \qquad \text{for } \kappa \in \mathit{Val}_\bot \tag{43}$$

$$[\![\ominus\, e]\!]_{\kappa} \mathrel{\widehat{=}} \bigsqcap\{\kappa' \in \mathit{Val}_\bot \mid \kappa = \mathit{eval}(\ominus, \kappa') \cdot [\![e]\!]_{\kappa'}\} \tag{44}$$

$$[\![e_1 \;\oplus\; e_2]\!]_{\kappa} \mathrel{\widehat{=}} \bigsqcap\{\kappa_1, \kappa_2 \in \mathit{Val}_\bot \mid \kappa = \mathit{eval}(\oplus, \kappa_1, \kappa_2) \cdot ([\![e_1]\!]_{\kappa_1} \parallel [\![e_2]\!]_{\kappa_2})\} \tag{45}$$

The evaluation $[\![\kappa_1]\!]_{\kappa_2}$ of a constant κ_1 to another constant κ_2 is represented by the command **idle**, which allows arbitrary interference and finite stuttering steps, followed by a test that the constants are equal (42). Note that the refinement $[\![\kappa_1]\!]_{\kappa_1} \sqsubseteq [\![\kappa_1]\!]_{\kappa_2}$ holds and hence $\bigsqcap\{\kappa_2 \in \mathit{Val}_\bot \cdot [\![\kappa_1]\!]_{\kappa_2}\} = [\![\kappa_1]\!]_{\kappa_1} = \mathbf{idle}$. Also, because constants κ_1

and κ_2 are independent of the state, the test $\boldsymbol{\tau}(\{\sigma \in \Sigma \mid \kappa_1 = \kappa_2\})$ is not subject to interference from the environment, and so $[\![\kappa_1]\!]_{\kappa_2} = \mathbf{idle}\,\boldsymbol{\tau}(\{\sigma \in \Sigma \mid \kappa_1 = \kappa_2\})\,\mathbf{idle}$.

To allow an assignment such as $x := x$ to be optimised to **skip** (which has no program steps), the evaluation of the expression x (and the assignment) must allow traces with no program steps. The core of the definition of $[\![x]\!]_\kappa$ is therefore the instantaneous test $\boldsymbol{\tau}(\{\sigma \in \Sigma \mid \sigma(x) = \kappa\})$ which is feasible only in states σ such that $\sigma(x) = \kappa$. A single reference to x is assumed to be atomic; if x was a structured variable with non-atomic access, a more complex definition would be needed.

Evaluation of a binary expression $e_1 \oplus e_2$ is nondeterministic because its value depends on the possibly changing values of variables in e_1 and e_2 (and similarly for unary expressions). For example, the traces of $[\![x+y]\!]_\kappa$ consist of all possible evaluations of both $[\![x]\!]_{\kappa_1}$ and $[\![y]\!]_{\kappa_2}$ such that $\kappa = \kappa_1 + \kappa_2$, which is represented as a nondeterministic choice.

$$[\![x+y]\!]_\kappa = \bigsqcap\{\kappa_1, \kappa_2 \mid \kappa = \kappa_1 + \kappa_2 \cdot [\![x]\!]_{\kappa_1} \parallel [\![y]\!]_{\kappa_2}\}$$

Subexpressions are evaluated in parallel to represent arbitrary order of subexpression evaluation (including parallel evaluation). For example, the evaluation of a (non-conditional) "and" of b_1 and b_2 to true corresponds to evaluating both b_1 and b_2 to true in parallel, i.e. $[\![b_1 \wedge b_2]\!]_{\mathsf{true}} = [\![b_1]\!]_{\mathsf{true}} \parallel [\![b_2]\!]_{\mathsf{true}}$.

$$\begin{aligned}
&[\![b_1 \wedge b_2]\!]_{\mathsf{true}} \\
&= \bigsqcap\{\kappa_1, \kappa_2 \mid \mathsf{true} = \mathit{eval}(\wedge, \kappa_1, \kappa_2) \cdot [\![b_1]\!]_{\kappa_1} \parallel [\![b_2]\!]_{\kappa_2}\} \\
&= \bigsqcap\{\kappa_1, \kappa_2 \mid \kappa_1 = \mathsf{true} \wedge \kappa_2 = \mathsf{true} \cdot [\![b_1]\!]_{\kappa_1} \parallel [\![b_2]\!]_{\kappa_2}\} \\
&= [\![b_1]\!]_{\mathsf{true}} \parallel [\![b_2]\!]_{\mathsf{true}}
\end{aligned}$$

For undefined expressions, such as a division by zero, the result of the evaluation is the undefined value $\bot_v$, for example, $[\![1/0]\!]_{\bot_v}$ succeeds.

5.2. Imperative programming constructs

An assignment, $x := e$, evaluates e to a value κ and then updates x to κ (46). While evaluation of the expression e is non-atomic, we assume that the act of storing of the value in the variable x is atomic. If the evaluation is undefined, the assignment aborts. If the test in a conditional (47) evaluates to true, c_1 is executed, if it evaluates to false, c_2 is executed, and if its evaluation is undefined it aborts. The body of a while loop (48) is repeated while its test evaluates to true. The additional $(\boldsymbol{\pi}(\mathsf{id})\,;\,\mathbf{idle})$ ensures it takes at least one program step on each iteration. If its test evaluates to false, it terminates, otherwise, if its test evaluation is undefined, it aborts.

$$x := e \mathrel{\widehat{=}} \bigsqcap\{\kappa \in \mathit{Val} \cdot [\![e]\!]_\kappa\, \mathit{update}(x, \kappa)\, \mathbf{idle}\} \sqcap ([\![e]\!]_{\bot_v}\, \bot_{\boldsymbol{\pi}}) \tag{46}$$

$$\mathbf{if}\, b\, \mathbf{then}\, c_1\, \mathbf{else}\, c_2 \mathrel{\widehat{=}} ([\![b]\!]_{\mathsf{true}}\, c_1) \sqcap ([\![b]\!]_{\mathsf{false}}\, c_2) \sqcap ([\![b]\!]_{\bot_v}\, \bot_{\boldsymbol{\pi}}) \tag{47}$$

$$\mathbf{while}\, b\, \mathbf{do}\, c \mathrel{\widehat{=}} ([\![b]\!]_{\mathsf{true}}\, c\,;\, \boldsymbol{\pi}(\mathsf{id})\,;\, \mathbf{idle})^\omega\, ([\![b]\!]_{\mathsf{false}} \sqcap ([\![b]\!]_{\bot_v}\, \bot_{\boldsymbol{\pi}})) \tag{48}$$

For example, the expansion of a simple assignment gives the following. For brevity we assume κ ranges over natural numbers only, but the full case does not change the reasoning.

$$\begin{aligned}
& x := 1 \\
= \;& \text{assignment definition} \\
& \bigsqcap\{\kappa \in \mathit{Val} \cdot [\![1]\!]_\kappa\, \mathit{update}(x, \kappa)\, \mathbf{idle}\} \sqcap ([\![1]\!]_{\bot_v}\, \bot_{\boldsymbol{\pi}}) \\
= \;& \text{split nondeterministic choice} \\
& ([\![1]\!]_1\, \mathit{update}(x, 1)\, \mathbf{idle}) \sqcap (\bigsqcap\{\kappa \in \mathit{Val} \mid \kappa \neq 1 \cdot [\![1]\!]_\kappa\, \mathit{update}(x, \kappa)\, \mathbf{idle}\}) \sqcap ([\![1]\!]_{\bot_v}\, \bot_{\boldsymbol{\pi}}) \\
= \;& \text{expand and simplify expression evaluation, } \top \text{ is left annihilator} \\
& (\mathbf{idle}\, \mathit{update}(x, 1)\, \mathbf{idle}) \sqcap (\bigsqcap\{\kappa \in \mathit{Val} \mid \kappa \neq 1 \cdot \mathbf{idle}\, \top\}) \sqcap (\mathbf{idle}\, \top) \\
= \;& \text{monotonicity of sequential composition, } \top \text{ is the greatest element} \\
& \mathbf{idle}\, \mathit{update}(x, 1)\, \mathbf{idle}
\end{aligned}$$

5.3. Termination and fairness

In the context of concurrency, termination of even the most obviously terminating program, such as "$x := 0$", relies on the program being scheduled long enough to complete. If the program is interrupted by its environment forever, it

does not terminate. For example, the program

$$x := 1 \,;\, ((\textbf{while}\, x \neq 0\, \textbf{do skip}) \parallel x := 0)$$

will not terminate unless the right-hand operand in the parallel composition is given a chance to set x to zero.

The command **term** is defined as $(\pi \sqcap \epsilon_\perp)^\star\, \epsilon_\perp^\omega$ which, although it only allows a finite number of program steps, does not preclude its environment taking over forever. Hence **term** $\sqsubseteq c$ is a slightly weak form of termination because it only requires termination of c provided its environment doesn't interrupt it forever.

Fair execution of a program rules out its environment interrupting it forever. The program that allows all behaviours except being interrupted forever can be expressed as

$$\textbf{fair} \mathrel{\widehat{=}} \epsilon_\perp^\star\, (\pi\, \epsilon_\perp^\star)^\omega$$

and fair execution of a command c can be represented by **fair** $\Cap c$, where weak conjunction ensures any aborting behaviour of c is preserved. A stronger form of termination can be expressed as

$$\textbf{fairterm} \mathrel{\widehat{=}} (\pi \sqcap \epsilon_\perp)^\star \;=\; \epsilon_\perp^\star\, (\pi\, \epsilon_\perp^\star)^\star$$

and requires a finite number of both program and environment steps. Interestingly,

$$\textbf{fairterm} = \textbf{fair} \Cap \textbf{term}$$

and hence if **term** $\sqsubseteq c$, by monotonicity, **fair** $\Cap$ **term** $\sqsubseteq$ **fair** $\Cap c$, that is, **fairterm** $\sqsubseteq$ **fair** $\Cap c$, and hence weak termination of c ensures strong termination if the execution of c is fair.

5.4. End-to-end specifications

In the rely-guarantee paradigm postconditions are end-to-end and allow only a finite number of program steps [22, 23]. A specification command, $(\!|\, q \,|\!)$, satisfies a binary relation q between its initial and final states [26].

$$(\!|\, q \,|\!) \mathrel{\widehat{=}} \bigsqcap \{\sigma \in \Sigma \cdot (\tau(\{\sigma\})\, \textbf{term}\, \tau(\{\sigma' \mid (\sigma, \sigma') \in q\}))\} \tag{49}$$

It may take any sequence of steps that collectively establish q. Sequences of steps that do not satisfy q are infeasible because of the final τ step. Note that this specification construct contains all possible implementations. In the literature it is often the case that q would initially be an atomic step, and this atomicity is then successively "split" into an implementation [3]. Brookes [5] and Dingel [9] use a notion of program equivalence modulo finite stuttering and mumbling.[7] The set of traces of a specification command is closed under finite stuttering and mumbling, and hence we can safely use the stronger notion of program equivalence as equality on sets of traces. For example, it is straightforward to show that a specification is closed under finite stuttering because $(\!|\, q \,|\!) \parallel \textbf{idle} = (\!|\, q \,|\!)$. Furthermore, (49) allows a trace with no program steps for initial states σ, such that $(\sigma, \sigma) \in q$, which isn't catered for by mumbling equivalence, however, it would be straightforward to generalise mumbling equivalence to allow this.

End-to-end specifications as standalone commands are in general unimplementable. The following program specifies the push of a value v onto a stack s. Here relations are expressed in predicative style, in which before state values of a variable s are represented by s and final state values by s' (as in Z [28, 16, 30] or TLA [25]). We let $[v]$ represent the singleton sequence containing v.

$$(\!|\, s' = [v] \frown s \,|\!)$$

There are no refinements of this specification to code because code does not constrain the environment to prevent modifications to s after it finishes its program steps. To refine this specification to code requires an assumption about the environment, in particular, that the environment does not modify s. In addition, this specification says nothing about which (other) variables may be modified. These parts of rely-guarantee reasoning are introduced in the next section.

[7] A trace t_1 is mumbling equivalent to a trace t_2 if t_1 contains a sub-trace $(\sigma_i\, \Pi\, \sigma_{i+1}\, \Pi\, \sigma_{i+2})$ and t_2 is the same as t_1 but with the sub-trace replaced by $(\sigma_i\, \Pi\, \sigma_{i+2})$, i.e. two consecutive program steps can be replaced by a single step with the same overall effect.

5.5. Rely and guarantee constructs

Because the model explicitly includes both environment and program steps, defining specification constructs to handle rely and guarantee conditions is straightforward.

5.5.1. Restricting the program: guarantees

The command $(\pi\text{-}\mathbf{restrict}\, g)$ restricts its program steps to satisfy the relation g and leaves its environment steps unconstrained. When combined with a command c using weak conjunction, $(\pi\text{-}\mathbf{restrict}\, g) \Cap c$, it allows only program steps of c that respect g (but does not prevent the combination from aborting if c can abort after a sequence of program steps satisfying g). This command encodes the *guarantee* condition of Jones [23].

$$(\pi\text{-}\mathbf{restrict}\, g) \mathrel{\widehat{=}} (\boldsymbol{\pi}(g) \sqcap \boldsymbol{\epsilon}_\bot)^\omega \tag{50}$$

Guarantees satisfy useful properties, such as the following.

$$(\pi\text{-}\mathbf{restrict}\, g_1) \Cap (\pi\text{-}\mathbf{restrict}\, g_2) = (\pi\text{-}\mathbf{restrict}\, g_1 \cap g_2)$$
$$(\pi\text{-}\mathbf{restrict}\, g_1) \parallel (\pi\text{-}\mathbf{restrict}\, g_2) = (\pi\text{-}\mathbf{restrict}\, g_1 \cup g_2)$$

A frame x on a command c, written $x{:}\, c$, constrains c to only modify the set of variables x; it can be defined as a guarantee to never modify variables outside x. The relation $\mathsf{id}(x)$ is the identity relation on the set of variables x (i.e. the relation $\{(\sigma, \sigma') \mid x \lhd \sigma = x \lhd \sigma'\}$, where $x \lhd \sigma$ is the state σ with its domain restricted to x) and $\overline{x}$ is the complement of the set of variables x.

$$x{:}\, c \mathrel{\widehat{=}} (\pi\text{-}\mathbf{restrict}\, \mathsf{id}(\overline{x})) \Cap c \tag{51}$$

As an example, recall pushing a value v onto a stack s. Typically an implementation is expected to modify s and only s.

$$(\!| s' = [v] \frown s |\!) \Cap (\pi\text{-}\mathbf{restrict}\, \mathsf{id}(\overline{s}))$$

which is equal to the abbreviation $s{:}\,(\!| s' = [v] \frown s |\!)$. Now any implementation of this specification may modify s only (in addition to any local variables, described below). However, it is still infeasible for the same reasons that a stand-alone end-to-end specification is infeasible; this problem will be rectified by the introduction of assumptions about the environment, described below.

Note that the program $s{:}\,(\!| s' = [v] \frown s |\!)$, is not the same as $(\!| s' = [v] \frown s \wedge \mathsf{id}(\overline{s}) |\!)$, because the latter may have intermediate program steps that modify variables other than s. Also, the latter specification constrains the environment steps in a way that the first command does not. In the first command, the environment may arbitrarily modify variables other than s, allowing the program to terminate in a state where variables other than s have been changed by the environment. For the second program, every trace is constrained to ensure that all variables other than s have a final value equal to their initial value, although it makes no constraints on any of the intermediate values of these variables. This difference is one of the most prominent between specification/refinement in sequential and parallel contexts.

5.5.2. Assumptions about the environment: rely

A command $(\epsilon\text{-}\mathbf{restrict}\, r)$ restricts every environment step to satisfy the relation r or abort, while allowing any program steps. A command $(\epsilon\text{-}\mathbf{rely}\, r)$ aborts if the environment does not respect r. The commands $(\epsilon\text{-}\mathbf{restrict}\, r)$ and $(\epsilon\text{-}\mathbf{rely}\, r)$ for environment steps perform similar roles in the language to those of guards and assertions on (pre-)states in the sequential refinement calculus.

The intention is for $(\epsilon\text{-}\mathbf{rely}\, r)$ to be composed with other specification commands using weak conjunction: any d such that $c \Cap (\epsilon\text{-}\mathbf{rely}\, r) \sqsubseteq d$ need refine c only when the environment respects r, as under any other circumstance the behaviour of d is unconstrained. This command encodes the *rely* condition of Jones [23].

$$(\epsilon\text{-}\mathbf{restrict}\, r) \mathrel{\widehat{=}} (\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot(r))^\omega \tag{52}$$
$$(\epsilon\text{-}\mathbf{rely}\, r) \mathrel{\widehat{=}} (\epsilon\text{-}\mathbf{restrict}\, r)\,(\boldsymbol{\tau} \sqcap \boldsymbol{\epsilon}(\overline{r})\, \bot_{\boldsymbol{\pi}}) \tag{53}$$

Environment guards and assumptions satisfy the following properties.

$$(\epsilon\text{-}\mathbf{restrict}\, r_1) \Cap (\epsilon\text{-}\mathbf{restrict}\, r_2) = (\epsilon\text{-}\mathbf{restrict}\, r_1 \cap r_2)$$
$$(\epsilon\text{-}\mathbf{rely}\, r_1) \Cap (\epsilon\text{-}\mathbf{rely}\, r_2) = (\epsilon\text{-}\mathbf{rely}\, r_1 \cap r_2)$$

More concretely, assume command p is a sequence of two atomic steps that initialises x to 0 and then increments x, for instance, in a language with atomic initialisation and increments, $x := 0 \,;\, x{+}{+}$. The traces of this command include all terminated and non-terminated traces with exactly two program steps interleaved with arbitrary environment steps:

$$p \;=\; \boldsymbol{\epsilon}_\bot^\omega\, \boldsymbol{\pi}(x' = 0)\, \boldsymbol{\epsilon}_\bot^\omega\, \boldsymbol{\pi}(x' = x+1)\, \boldsymbol{\epsilon}_\bot^\omega \;.$$

In general we may not state anything about the final value of x in the terminating traces of p since the environment steps may modify x. If we guard the environment to restrict to steps that do not modify x, however, we obtain the following types of traces.

$$p \Cap (\epsilon\text{-}\mathbf{restrict}\, x' = x) \;=\; \boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = 0)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = x+1)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega$$

In this new program, every environment step leaves x unmodified or aborts, and so we can deduce that x equals 1 in the final state of any terminating behaviour.

In practice, restricting the environment is too strong for specifications. The use of a rely command, however, allows all possible behaviours of the environment.

$$\begin{aligned} p \Cap (\epsilon\text{-}\mathbf{rely}\, x' = x) \;=\; & (\boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = 0)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = x+1)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega) \sqcap \\ & (\boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\epsilon}_\bot(x' \neq x)\, \bot_{\boldsymbol{\pi}}) \sqcap \\ & (\boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = 0)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\epsilon}_\bot(x' \neq x)\, \bot_{\boldsymbol{\pi}}) \sqcap \\ & (\boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = 0)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\pi}(x' = x+1)\, \boldsymbol{\epsilon}_\bot(x' = x)^\omega\, \boldsymbol{\epsilon}_\bot(x' \neq x)\, \bot_{\boldsymbol{\pi}}) \end{aligned}$$

The traces of $p \Cap (\epsilon\text{-}\mathbf{rely}\, x' = x)$ therefore contain the traces where the environment does not modify x, establishing x equals 1 in traces that terminate, as well as traces where the environment does not respect the rely condition, at which point any behaviour is possible. When used in a specification, the rely command therefore absolves the implementation of establishing anything when the condition is not respected by its environment; similarly if a sequential program begins execution in a state that does not satisfy its precondition it need not establish the postcondition.

We can complete the specification of push onto a stack in an environment that does not modify s by adding the appropriate rely.

$$s{:}\,(\!|\, s' = [v] \frown s \,|\!) \Cap (\epsilon\text{-}\mathbf{rely}\, \mathsf{id}(s)) \tag{54}$$

This is now feasible: it may be implemented by the typical two-line linked list code (after data refining the representation). Traces of the implementation in which the stack is not modified by its environment must satisfy the end-to-end relation $s' = [v] \frown s$, while traces of the implementation in which the stack is modified by the environment are always legal, since the specification aborts under such circumstances.

The rely $\mathsf{id}(s)$ is a strong assumption, essentially requiring that s is local to the program (although its value may be read by the environment). It is often the case that the program and environment communicate via shared variables, in which case a weaker rely is used. For instance, if i is a shared index into an array for concurrent search, an implementation may allow each program to assume that i only increases, never decreases (given by the command $(\epsilon\text{-}\mathbf{rely}\, i' \geq i)$). The weakest assumption is $(\epsilon\text{-}\mathbf{rely}\, \varnothing)$, making no assumptions about what the environment may modify. This is usually too weak for feasible implementations, for instance, $(\!|\, q \,|\!) \Cap (\epsilon\text{-}\mathbf{rely}\, \varnothing) = (\!|\, q \,|\!)$. More fully developed refinements in a rely-guarantee context appear in [24, Sect 2.3] and [14, Sect. 11].

Traditionally rely-guarantee reasoning is formulated as a quintuple statement $\{p, r\}c\{g, q\}$, which states that, assuming that the initial state satisfies p and each step of the environment satisfies r, then c terminates and establishes q and furthermore every program step of c satisfies the guarantee g [23]. A rely-guarantee quintuple can be rewritten as a refinement statement in our language as follows.

$$\{p\}(\!|\, q \,|\!) \Cap (\epsilon\text{-}\mathbf{rely}\, r) \Cap (\pi\text{-}\mathbf{restrict}\, g) \sqsubseteq c \tag{55}$$

Note the separation of concerns into pre, post, rely and guarantee conditions allowing separate reasoning about each component.

5.5.3. Further properties of environments

Because environment guards and relies are novel to our approach we explore some useful properties of these commands.

Lemma 1 (eguard-env). For a binary relation r, $\quad (\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{rely}\, r) = (\epsilon\text{-}\mathbf{restrict}\, r)$.

Proof. The proof expands the definition of $(\epsilon\text{-}\mathbf{rely}\, r)$ and simplifies.

$$
\begin{array}{ll}
& (\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{rely}\, r) \\
= & (\epsilon\text{-}\mathbf{restrict}\, r) \Cap ((\epsilon\text{-}\mathbf{restrict}\, r) \sqcap (\epsilon\text{-}\mathbf{restrict}\, r)\, \boldsymbol{\epsilon}(\bar{r})\, \bot_{\boldsymbol{\pi}}) \\
= & ((\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{restrict}\, r)) \sqcap ((\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{restrict}\, r)\, \boldsymbol{\epsilon}(\bar{r})\, \bot_{\boldsymbol{\pi}}) \\
= & (\epsilon\text{-}\mathbf{restrict}\, r) \sqcap ((\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{restrict}\, r)\, \boldsymbol{\epsilon}(\bar{r})\, \bot_{\boldsymbol{\pi}}) \\
= & (\epsilon\text{-}\mathbf{restrict}\, r)
\end{array}
$$

For the last step one needs to show $(\epsilon\text{-}\mathbf{restrict}\, r) \sqsubseteq ((\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{restrict}\, r)\, \boldsymbol{\epsilon}(\bar{r})\, \bot_{\boldsymbol{\pi}})$, which can be shown by induction. □

Lemma 2. For a binary relation r, and commands c and d,

$$(\epsilon\text{-}\mathbf{restrict}\, r) \Cap c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d \quad \Leftrightarrow \quad c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d\ .$$

Proof. Assuming the left side, the right follows by the reasoning below because $\mathbf{chaos} \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r)$.

$$c \;=\; \mathbf{chaos} \Cap c \;\sqsubseteq\; (\epsilon\text{-}\mathbf{restrict}\, r) \Cap c \;\sqsubseteq\; (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d$$

The proof from right to left is as follows.

$$
\begin{array}{ll}
& c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d \\
\Rightarrow & (\epsilon\text{-}\mathbf{restrict}\, r) \Cap c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d \\
\Leftrightarrow & (\epsilon\text{-}\mathbf{restrict}\, r) \Cap c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d
\end{array}
$$

□

The literature on rely-guarantee often makes use of a refinement operator $c \sqsubseteq_r d$ requiring d to refine c provided all environment steps satisfy r [6]. More formally, the set of traces of d restricted to those with all environment steps satisfying r is contained in the traces of c. The set of traces of d restricted to those with all environment steps satisfying r corresponds to $(\epsilon\text{-}\mathbf{restrict}\, r) \Cap d$, and hence the refinement $c \sqsubseteq_r d$ is equivalent to

$$c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d\ . \tag{56}$$

Hence by making use of a richer set of commands, one can avoid the need to introduce the family of refinement relations "$\sqsubseteq_r$" indexed by r. Property (56) is subtly different to the refinement

$$(\epsilon\text{-}\mathbf{rely}\, r) \Cap c \sqsubseteq d \tag{57}$$

but only for commands c that restrict environment steps. In general, (57) implies (56) but not vice versa. To see that the reverse does not hold take c to be $(\epsilon\text{-}\mathbf{restrict}\, r)$ and d to be $\mathbf{chaos}$. Refinement (56) reduces to $(\epsilon\text{-}\mathbf{restrict}\, r) \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap \mathbf{chaos} = (\epsilon\text{-}\mathbf{restrict}\, r)$, which holds, but (57) reduces to $(\epsilon\text{-}\mathbf{rely}\, r) \Cap (\epsilon\text{-}\mathbf{restrict}\, r) \sqsubseteq \mathbf{chaos}$, which does not hold in general because the left side is equivalent to $(\epsilon\text{-}\mathbf{restrict}\, r)$ by Lemma 1.

Theorem 3. For a binary relation r, and commands c and d, $(\epsilon\text{-}\mathbf{rely}\, r) \Cap c \sqsubseteq d \;\Rightarrow\; c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d$.

Proof.

$$
\begin{array}{ll}
& (\epsilon\text{-}\mathbf{rely}\, r) \Cap c \sqsubseteq d \\
\Rightarrow & \text{as weak conjunction is monotone} \\
& (\epsilon\text{-}\mathbf{restrict}\, r) \Cap (\epsilon\text{-}\mathbf{rely}\, r) \Cap c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d \\
\Leftrightarrow & \text{by Lemma 1} \\
& (\epsilon\text{-}\mathbf{restrict}\, r) \Cap c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d \\
\Leftrightarrow & \text{by Lemma 2} \\
& c \sqsubseteq (\epsilon\text{-}\mathbf{restrict}\, r) \Cap d
\end{array}
$$

□

This law illustrates the connection between environment restriction and relies with respect to refinement, much like the connection between guards and assertions for preconditions from the sequential refinement calculus. Given (normal) programs c and d that do not restrict the environment, making a rely assumption in the specification c is equivalent to restricting the environment in d when showing refinement.

5.6. Local variable blocks

In the context of interference from concurrent programs, a local variable cannot be affected by any programs external to the block in which it is declared and the body of the local variable block cannot reference[8] (and hence modify) any more global variable with the same name. The definition of a local variable block can be built from the language primitives defined earlier, including the operator $c \setminus x$, which has traces that are the same as c except that x is unrestricted (27).

The command $(\mathbf{local}\, x \cdot c)$ localises x within c and hence it restricts all environment steps of c to not modify x, i.e. they satisfy the relation $\mathsf{id}(x)$. The program steps of the local variable block do not change the value of a non-local occurrence of the variable x, i.e. $(\pi\text{-}\mathbf{restrict}\, \mathsf{id}(x))$. For a local variable block $(\mathbf{vardecl}\, x \cdot c)$, c is run in an environment in which x is local. To allow for allocation and deallocation of x, the block may idle before and after execution (implicitly such idle steps may modify the local x).

$$\mathbf{local}\, x \cdot c \mathrel{\widehat{=}} (\pi\text{-}\mathbf{restrict}\, \mathsf{id}(x)) \Cap (c \Cap (\epsilon\text{-}\mathbf{restrict}\, \mathsf{id}(x))) \setminus x$$
$$\mathbf{vardecl}\, x \cdot c \mathrel{\widehat{=}} \mathbf{idle}\, (\mathbf{local}\, x \cdot c)\, \mathbf{idle}$$

From this definition it follows that $(\mathbf{vardecl}\, x \cdot c) = (\pi\text{-}\mathbf{restrict}\, \mathsf{id}(x)) \Cap (\mathbf{vardecl}\, x \cdot c)$ and any refinement of c within the block can assume there is no interference on x. If the body of a local variable block guarantees the relation g on every program step, the block itself respects both $\mathsf{id}(x)$ and $g \setminus x$, where $(\sigma, \sigma') \in (g \setminus x) \Leftrightarrow (\exists\, \sigma_1, \sigma_1' \cdot x \lhd \sigma = x \lhd \sigma_1 \wedge x \lhd \sigma' = x \lhd \sigma_1' \wedge (\sigma_1, \sigma_1') \in g)$.

$$(\mathbf{vardecl}\, x \cdot (\pi\text{-}\mathbf{restrict}\, g) \Cap c) = (\pi\text{-}\mathbf{restrict}\, g \setminus x \cap \mathsf{id}(x)) \Cap (\mathbf{vardecl}\, x \cdot (\pi\text{-}\mathbf{restrict}\, g) \Cap c)$$

An assumption about environment steps satisfying r can be propagated into a local variable block by ignoring the global variable x (if any) and assuming the local x cannot be changed by the environment.

$$(\epsilon\text{-}\mathbf{rely}\, r) \Cap (\mathbf{vardecl}\, x \cdot c) \sqsubseteq (\mathbf{vardecl}\, x \cdot (\epsilon\text{-}\mathbf{rely}(r \setminus x) \cap \mathsf{id}(x)) \Cap c)$$

Consider the variable declaration command $d \mathrel{\widehat{=}} (\mathbf{vardecl}\, x \cdot c)$ in an environment e, i.e. $d \parallel e$. The intention is that c may modify x, but that e may not modify d's x. The former intention is implicitly allowed by d. The latter constraint is enforced by the environment guard $\mathsf{id}(x)$ in $(\mathbf{local}\, x \cdot c)$. Typical approaches to local variables require a syntactic constraint on e that it never modify x, but this is problematic if e declares its own copy of x, for instance, if d and e are instances of the same command. In that case a further implicit or explicit renaming scheme is introduced to remove the ambiguity. Such schemes create significant overhead if enforced fully (for instance, environments in Plotkin-style semantics). The approach taken here avoids such syntactic concerns by "hiding" c's modifications of x from e. This is given by $c \setminus x$ in $(\mathbf{local}\, x \cdot c)$, which removes the modifications of x in the program steps of c. These modifications are then replaced by a guarantee that d does not modify x, i.e. $(\pi\text{-}\mathbf{restrict}\, \mathsf{id}(x))$, as expected: if x *is* (re)declared in e, d does not modify it, since all references to x in d are local.

More concretely, the traces of $(\mathbf{vardecl}\, x \cdot c_1) \parallel (\mathbf{vardecl}\, x \cdot c_2)$ do not interfere with each other on x. The declaration of x as local to each program is a semantic restriction, rather than a syntactic restriction.

5.7. Atomic abstract specifications and linearisability

We have previously discussed end-to-end specifications, which allow implementations that comprise multiple steps that modify (potentially shared) data. In concurrent contexts in which many programs modify some central shared data, end-to-end specifications are not appropriate: either a multi-stage update to shared data may interfere with the other programs, or other programs may modify the shared data after completing an operation, falsifying the end-to-end requirement. For situations where multiple programs are expected to concurrently make calls that access/update shared data, some form of locking or other mechanism is required to ensure the data remains consistent. One of the most widely known correctness criterion for shared data with atomic actions is linearisability [19]. We demonstrate below how the wide-spectrum language is applicable in this widely-used context.

To describe atomic execution it useful to define an abstract *atomic step* command $\langle r \rangle$, which performs exactly one atomic step that satisfies r, in addition to which it may idle before and after.

$$\langle r \rangle \mathrel{\widehat{=}} \mathbf{idle}\, \boldsymbol{\pi}(r)\, \mathbf{idle} \tag{58}$$

[8] We assume there is no aliasing of variable names otherwise restrictions on aliasing need to be included.

5.7.1. Concurrent objects

The concept of a concurrent object is shared data x which may be modified by programs (indexed over the finite set P) performing operations (drawn from set OP) concurrently. This can be abstractly stated in the wide-spectrum language using local variables and generalising binary parallel composition.

$$\mathbf{vardecl}\, x \cdot \parallel_{p \in P} \left(\prod_{op \in OP} op \right)^{\omega} \tag{59}$$

Each program repeatedly executes operations on x. The interleaving of multiple modifications to s results in undesirable interference. Traditionally such interference is avoided through the introduction of locks, although this approach has drawbacks associated with deadlock and unreliable processors. Another approach is to use atomic hardware primitives such as compare-and-swap (CAS) and retry loops to avoid deadlock. In either approach, the intuition is that the operations appear to take effect atomically from the perspective of other programs. This notion of correctness is formalised by Herlihy & Wing's definition of *linearisability*. Although the original definition is in terms of histories of operation invocation and response events (rather than traces of states), it is often equivalently recast in terms of atomic specifications on abstract data, which are data refined to (usually pointer-based) implementations in code.

5.7.2. Atomic operations on concurrent objects

An operation to perform an atomic step satisfying the relation q on a data structure x has the following form, which includes an assumption p of the initial state and a rely r of the environment.

$$\{p\}\ x{:}\, \langle q \rangle \Cap (\epsilon\text{-}\mathbf{rely}\, r) \tag{60}$$

Note that finite stuttering is allowed in addition to the single q-step. Any guarantee condition can be included as a conjunct within q. As an example, below we define *push* and *pop* operations on an abstract stack represented as a sequence s. Recall that sequence concatenation is represented by the operator "$\frown$" and $[v]$ represents the singleton sequence containing v.

$$push(s, v) \mathrel{\widehat{=}} s{:}\, \langle s' = [v] \frown s \rangle \Cap (\epsilon\text{-}\mathbf{rely}\, v' = v) \tag{61}$$

$$pop(s, v) \mathrel{\widehat{=}} s, v{:}\, \langle s = [v'] \frown s' \vee (s = s' = [\,] \wedge v' = \mathsf{none}) \rangle \Cap (\epsilon\text{-}\mathbf{rely}\, v' = v) \tag{62}$$

A call $push(s, v)$ atomically adds value v to the top of the stack. An implementation of $push(s, v)$ may make any finite number of stuttering steps before and after the main atomic step that modifies s; this is important for implementations that may use a different representation. The *push* operation need not assume anything about the initial state and assumes that the environment does not modify the parameter v (and that the type of s remains a stack, which we leave implicit). A call $pop(s, v)$ returns the popped value in v, or the special value none if the stack is empty at the time the atomic step of the pop occurs.

Encapsulating the stack as a concurrent object gives the following instance of (59).

$$Stack \mathrel{\widehat{=}} \mathbf{vardecl}\, s \cdot \parallel_{p \in P} (\mathbf{vardecl}\, v \cdot (push(s, v) \sqcap pop(s, v))^{\omega})$$

In this specification multiple programs concurrently modify s by atomically pushing to or popping from it, as recorded in the local variable v of each process.

Compare the specification of atomic push on a stack in (61) to that using an end-to-end specification in (54). In the former no assumption is made about s (except that implicitly it respects its type), while in the latter s is assumed not to be changed by the environment. The former is a weaker assumption about the environment, and hence more desirable in general, but with the trade-off that the push itself must be completed in exactly one program step. The specification in (54) allows multiple steps to achieve the specification since it may assume no interference on s. A further point of difference is that a guarantee is redundant for the atomic specification (61) since it takes only one step.

6. Progress properties

In addition to "safety" properties such as establishing an end-to-end relation, concurrent programs are often required to establish "progress" properties regarding overall behaviour. In Section 6.1 temporal logic is recast into our refinement framework, and then used in Section 6.2 to express the well-known "lock-free" property, and two related concepts.

6.1. Temporal logic

Linear temporal logic (LTL) [27] describes properties of infinite traces. It is straightforward to encode a temporal logic formula f as a command $[\![f]\!]$ such that a trace tr satisfies f if and only if $tr \in [\![f]\!]$ and hence a command c satisfies f if and only if $[\![f]\!] \sqsubseteq c$. As well as standard LTL formulae, one can include formulae representing a single program or environment step satisfying a relation r, for which the notations $\Pi(r)$ and $\mathrm{E}(r)$, respectively, are used. This makes it possible to devise temporal logic formulae that relate program and environment behaviour. In the following, p is a single-state predicate, r is a binary relation on states, h is a temporal logic formula, and x is a bound variable from a recursion.

LTL formula f	$[\![f]\!]$	LTL formula f	$[\![f]\!]$
p	$\boldsymbol{\tau}(p)\,;\,\bot_{\boldsymbol{\pi}}$	$\bigcirc h$	$(\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot)\,;\,[\![h]\!]$
$\Pi(r)$	$\boldsymbol{\pi}(r)\,;\,\bot_{\boldsymbol{\pi}}$	$\Diamond h$	$(\boldsymbol{\pi} \sqcap \boldsymbol{\epsilon}_\bot)^\star\,;\,[\![h]\!]$
$\mathrm{E}(r)$	$\boldsymbol{\epsilon}(r)\,;\,\bot_{\boldsymbol{\pi}}$	$\Box h$	$\mu x \cdot [\![h]\!] \sqcup \bigcirc x$
$h_1 \wedge h_2$	$[\![h_1]\!] \sqcup [\![h_2]\!]$	x	x
$h_1 \vee h_2$	$[\![h_1]\!] \sqcap [\![h_2]\!]$		

The encoding of the primitive single state LTL formula p requires p to hold in the initial state, after which any behaviour ($\bot_{\boldsymbol{\pi}}$) is possible. In addition to this familiar LTL formula, in our language we allow relations r to be LTL formulae, stating whether the program or environment must establish r. This distinction is used below to require that the environment always take some step. Conjunction and disjunction of formulae straightforwardly use intersection ($\sqcup$) and union ($\sqcap$) over sets of traces. The LTL formula for next-step $\bigcirc h$ requires exactly one step (program or environment) to occur and then h to be satisfied. The LTL formula for eventually-holds $\Diamond h$ requires h to hold after a finite number of steps. The encoding for always-holds ($\Box h$) uses a typical recursive encoding where h is required to hold in every step.

These definitions highlight the expressive power of a language based on primitives, and are exploited below to describe some common progress properties of concurrent programs. Further ways to develop a temporal logic for our wide spectrum language, including interval temporal logic, are outside the scope of this paper.

6.2. Non-blocking operations

As it is usually conceived, *lock-freedom* allows potential non-termination of an operation provided the system as a whole continues to make progress. That is, the non-terminated operation does not "lock" the system, and an infinite number of operations (executed by other parallel processes) may still complete. For a data structure s, for which the shared data is modified at most once per operation (successful or otherwise), system-wide progress can be detected by the data structure changing; more specifically in our context, by a step $\mathrm{E}(s' \neq s)$.

Related concepts include the stronger *wait-freedom*, that means every program completes regardless of interference, and the weaker *obstruction-freedom*, that states that programs need make progress only in the absence of other programs [18]. Given an atomic operation op that operates on data structure s, we can specify these properties straightforwardly using distinguished environment steps and LTL.

$$op_{of} \mathrel{\widehat{=}} op \sqcap ((\pi\text{-}\mathbf{restrict}\,\mathsf{id}) \Cap \Box\Diamond \mathrm{E}) \tag{63}$$

$$op_{lf} \mathrel{\widehat{=}} op \sqcap ((\pi\text{-}\mathbf{restrict}\,\mathsf{id}) \Cap \Box\Diamond \mathrm{E}(s' \neq s)) \tag{64}$$

$$op_{wf} \mathrel{\widehat{=}} op \tag{65}$$

An obstruction-free operation is weaker than a lock-free operation: in addition to successful completion of the operation, an implementation of op_{of} may fail to terminate if the environment continues to take steps (always eventually the trace contains an environment step). The restriction $(\pi\text{-}\mathbf{restrict}\,\mathsf{id})$ ensures that the implementation may not change the state in this case. The lock-free specification op_{lf} is similar, but an implementation may fail to terminate only if

the environment continually modifies s, indicating system-wide progress. A valid implementation of op_{lf} therefore includes a potentially infinite loop that repeats whenever interference is detected, and otherwise completes the operation. Wait-freedom requires termination regardless of the environment; a wait-free operation is the default case. This classification makes it clear that $op_{of} \sqsubseteq op_{lf} \sqsubseteq op_{wf}$, as expected. If the environment performs no steps, op_{of} implements op and if the environment does not change s, op_{lf} implements op.

$op \sqsubseteq op_{of} \Cap (\epsilon\text{-}\mathbf{restrict}\, \varnothing)$

$op \sqsubseteq op_{lf} \Cap (\epsilon\text{-}\mathbf{restrict}\, \mathsf{id}(s))$

This classification is simpler than many in the literature (though less general than Dongol [10]). We argue that this is achieved by using Aczel traces to specify the environment, and applying LTL where appropriate. Note that properties of guarantees and LTL developed in different contexts may be reused to tackle nonblocking programs.

7. Conclusions

Our goal in designing the semantic model presented in this paper is to provide an approach that supports a concurrent wide-spectrum language that includes both programming language constructs and specification constructs, including constructs that support rely and guarantee conditions. A novelty of our approach is that we make use of a set of primitive commands, including $\pi(r)$, $\epsilon(r)$ and $\tau(p)$, that are quite close to the atomic steps of the underlying semantics and provide an expressive framework for defining diverse concepts from concurrency theory. More complex commands are defined using these primitive commands and the primitive operators. The use of such basic primitives allows one to make fine distinctions in the semantics, e.g. we considered three "no-op" commands: τ, **skip**, and **idle**.

The semantics of the primitives is given in terms of Aczel traces [1], similar to the approach of De Roever [8] but extended to support aborting and infeasible specification constructs. Alternative approaches to specifying the core language are possible, for example using an operational semantics that distinguishes program and environment steps. As long as the properties of the primitives are maintained, the definition of the wide-spectrum language constructs in terms of the primitives requires no change.

Brookes gives a semantics for a parallel language in terms of *action traces* [4]. He gives a set of ten primitive *actions* that form the atomic units of a program's execution. Because he supports concurrent separation logic, his model of the state includes a store, a heap and a finite set of resources. Given such a model of the state, his primitive actions can all be defined in terms of a command of the form $\pi(r)$, in which the relation r determines the effect of the primitive. More generally, the primitives could be defined as a command of the form $opt(r)$, which allows the step to be elided for initial states in which the command has no effect on the state.

In [14, 24] a refinement calculus for developing programs in the rely-guarantee style is presented, using an underlying semantics similar to (but less general than) that presented here. Many of the laws developed in that paper for dealing with assumptions about initial states and the environment may be carried over to reasoning about those aspects of atomic specifications, which is an advantage of a unified concurrency framework.

Acknowledgements. This research was supported by ARC grant DP130102901. This paper has benefited from feedback from Julian Fell, Cliff Jones, Graeme Smith, Kim Solin and Kirsten Winter.

References

[1] P. H. G. Aczel. On an inference rule for parallel composition, 1983. Private communication to Cliff Jones `http://homepages.cs.ncl.ac.uk/cliff.jones/publications/MSs/PHGA-traces.pdf`.

[2] R.-J. R. Back and J. von Wright. *Refinement Calculus: A Systematic Introduction*. Springer, New York, 1998.

[3] R.J.R. Back. A method for refining atomicity in parallel algorithms. In Eddy Odijk, Martin Rem, and Jean-Claude Syre, editors, *PARLE '89 Parallel Architectures and Languages Europe*, volume 366 of *LNCS*, pages 199–216. 1989.

[4] S. Brookes. A semantics for concurrent separation logic. *Theoretical Computer Science*, 375(13):227–270, 2007.

[5] S. D. Brookes. Full abstraction for a shared-variable parallel language. *Information and Computation*, 127(2):145–163, June 1996.

[6] J. W. Coleman and C. B. Jones. A structural proof of the soundness of rely/guarantee rules. *Journal of Logic and Computation*, 17(4):807–841, 2007.

[7] F.S. de Boer, U. Hannemann, and W.-P. de Roever. Formal justification of the rely-guarantee paradigm for shared-variable concurrency: a semantic approach. In Jeannette Wing, Jim Woodcock, and Jim Davies, editors, *FM99 Formal Methods*, volume 1709 of *Lecture Notes in Computer Science*, pages 1245–1265. Springer Berlin / Heidelberg, 1999.

[8] W.-P. de Roever. *Concurrency Verification: Introduction to Compositional and Noncompositional Methods*. Cambridge University Press, 2001.

[9] J. Dingel. A refinement calculus for shared-variable parallel and distributed programming. *Formal Aspects of Computing*, 14(2):123–197, 2002.

[10] Brijesh Dongol. Formalising progress properties of non-blocking programs. In Zhiming Liu and Jifeng He, editors, *Formal Methods and Software Engineering*, volume 4260 of *Lecture Notes in Computer Science*, pages 284–303. Springer Berlin Heidelberg, 2006.

[11] R. W. Floyd. Assigning meaning to programs. *Math. Aspects of Comput. Sci.*, 19:19–32, 1967.

[12] Simon Foster, Frank Zeyda, and Jim Woodcock. Isabelle/UTP: A mechanised theory engineering framework. In David Naumann, editor, *Unifying Theories of Programming - 5th International Symposium, UTP 2014, Singapore, May 13, 2014, Revised Selected Papers*, volume 8963 of *Lecture Notes in Computer Science*, pages 21–41. Springer, 2014.

[13] I. J. Hayes. Generalised rely-guarantee concurrency: An algebraic foundation. *arXiv*, 1603.01776:1–23, 2015. `http://arxiv.org/abs/1603.01776`.

[14] I. J. Hayes, C. B. Jones, and R. J. Colvin. Laws and semantics for rely-guarantee refinement. Technical Report CS-TR-1425, Newcastle University, July 2014.

[15] I. J. Hayes and M. Utting. Deadlines are termination. In D. Gries and W.-P. de Roever, editors, *IFIP TC2/WG2.2, 2.3 International Conference on Programming Concepts and Methods (PROCOMET'98)*, pages 186–204. Chapman and Hall, 1998.

[16] Ian Hayes, editor. *Specification Case Studies*. Prentice Hall International, second edition, 1993.

[17] Ian J. Hayes and Larissa Meinicke. Invariants, well-founded statements and real-time program algebra. In Cliff. B. Jones, P. Pihlajasaari, and Jun Sun, editors, *Formal Methods (FM 2014)*, volume 8442 of *LNCS*, pages 318–334. Springer, 2014.

[18] M. Herlihy, V. Luchangco, and M. Moir. Obstruction-free synchronization: Double-ended queues as an example. In *ICDCS '03: Proceedings of the 23rd International Conference on Distributed Computing Systems*, pages 522–529, Washington, DC, USA, 2003. IEEE Computer Society.

[19] Maurice P. Herlihy and Jeannette M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463 – 492, 1990.

[20] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 583, October 1969.

[21] C. A. R. Hoare, B. Möller, G. Struth, and I. Wehrman. Concurrent Kleene Algebra and its foundations. *J. Log. Algebr. Program.*, 80(6):266–296, 2011.

[22] C. B. Jones. *Development Methods for Computer Programs including a Notion of Interference*. PhD thesis, Oxford University, June 1981. Available as: Oxford University Computing Laboratory (now Computer Science) Technical Monograph PRG-25.

[23] C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Transactions on Programming Languages and Systems*, 5(4):596–619, 1983.

[24] C. B. Jones, I. J. Hayes, and R. J. Colvin. Balancing expressiveness in formal approaches to concurrency. *Formal Aspects of Computing*, 27(3):475–497, May 2015.

[25] Leslie Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison Wesley, 2003.

[26] C. C. Morgan. The specification statement. *ACM Trans. Prog. Lang. and Sys.*, 10(3):403–419, July 1988.

[27] Amir Pnueli. The temporal logic of programs. In *Foundations of Computer Science, 18th Annual Symposium on*, pages 46–57. IEEE, 1977.

[28] J.M. Spivey. *The Z Notation: A Reference Manual*. Prentice Hall International, second edition, 1992.

[29] J. von Wright. Towards a refinement algebra. *Science of Computer Programming*, 51:23–45, 2004.

[30] Jim Woodcock and Jim Davies. *Using Z: Specification, Refinement, and Proof*. Prentice Hall International, 1996.